Creating And Running Apps On The Microsoft Azure Hybrid Cloud

Live Streaming With FFMPEG

₹ 100

Volume: 04 | Issue: 03 | Pages: 100 | December 2015

OpenSource For You

THE COMPLETE MAGAZINE ON OPEN SOURCE

An EFY GROUP Publication

FOSS Power For Windows!

Running VirtualBox On Windows

Create An Exe File For Python Applications In Windows

Daily-Use Free Software For Windows

Apache mahout

The Recommender System For Big Data

An Interview With Piyush Jha, AVP, Product Engineering, GlobalLogic

eNlight now becomes more stronger with US PATENT for

"Method and System For Real Time Detection of Resource Requirement and Automatic Adjustments"

esds
enabling futurability

Mumbai | Delhi | Bengaluru | Nashik | Leeds (UK)

ESDS SOFTWARE SOLUTION PVT. LTD.

Toll Free No. : 1800 209 3006 | Fax : +91 95 95 247 247
Email : relationship@esds.co.in | Website : www.esds.co.in

An EFYGROUP EVENT | Exhibition & Knowledge Partner ELCINA

www.efyexpo.com

INDIA ELECTRONICS WEEK

January 11th-13th, 2016 • Bengaluru

Bangalored!

In 2016, EFY Expo will be held at
Bangalore as part of India Electronics Week.
Namma Bengaluru!

Electronics ForYou expo

Innovate. Design. Manufacture. Source.

January 11th-13th, 2016 • Bengaluru

For more information, talk to us at +91-11-40596605 or email at growmybiz@efy.in

Gold Partner

KEYSIGHT TECHNOLOGIES | ARTS

Associate Partners

MOUSER ELECTRONICS | messung | erfi | Aplab | JUKI | emst | DRIVE

Tea/Coffee Partner

KWK RESISTORS INDIA PVT LTD.

Supporting Association Partners

CEAMA | ESC | SCCI | COWE | IESA | IPC INDIA | SMTA | TEMA | IEEE

International Partners

CEC | CEIEC | TAITRA

Community Partners

ICOO SALES | ICTBLR | DIY Hacking

Media Partners

Online

electronicsforu | EFYTimes | OpenSource | elsiEinder | EXPLARA | 10Times | ComponentTrade | LED Bazaar

Print

ELECTRONICS BAZAAR | LED | OpenSource | EB ELECTRONICS WORLD

Contents

ADMIN

DEVELOPERS

REGULAR FEATURES

An EFY GROUP EVENT

India's first mega event for all the stake-holders of Internet of Things ecosystem to come together.

iotshow.in

January 11th-13th, 2016

Bangalore International Exhibition Centre, Bengaluru

To increase your sales by exhibiting in this event, Call: **Jyotsna (011-40596605)** or write to us at: **efyexpo@efy.in**

Co-located at INDIA ELECTRONICS WEEK

January 11th-13th, 2016
Bengaluru

Contents

Editor
RAHUL CHOPRA

Editorial, Subscriptions & Advertising
DELHI (HQ)
D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020
Ph: (011) 26810602, 26810603; Fax: 26817563
E-mail: info@efy.in

Missing Issues
E-mail: support@efy.in

BENGALURU
Ph: (080) 25260394, 25260023
E-mail: efyblr@efy.in

Customer Care
E-mail: support@efy.in

Back Issues
Kits 'n' Spares
New Delhi 110020
Ph: (011) 26371661, 26371662
E-mail: info@kitsnspares.com

Advertising
CHENNAI
Ph: (044) 42994363
E-mail: efyenq@efy.in

HYDERABAD
Ph: (040) 67172633
E-mail: efyenq@efy.in

MUMBAI
Ph: (022) 24950047, 24928520
E-mail: efymum@efy.in

PUNE
Ph: 08800295610/ 09870682995
E-mail: efypune@efy.in

GUJARAT
Ph: (079) 61344948
E-mail: efyahd@efy.in

JAPAN
Tandem Inc., Ph: 81-3-3541-4166
E-mail: tandem@efy.in

SINGAPORE
Publicitas Singapore Pte Ltd
Ph: +65-6836 2272
E-mail: publicitas@efy.in

UNITED STATES
E & Tech Media
Ph: +1 860 536 6677
E-mail: veroniquelamarque@gmail.com

CHINA
Power Pioneer Group Inc.
Ph: (86 755) 83729797, (86) 13923802595
E-mail: powerpioneer@efy.in

TAIWAN
J.K. Media, Ph: 886-2-87726780 ext. 10
E-mail: jkmedia@efy.in

Exclusive News-stand Distributor (India)
IBH BOOKS AND MAGAZINES DISTRIBUTORS LTD
Unit No.10, Bezzola Complex , Opp. Suman Nagar,
Sion Trombay Road, Chembur,
Mumbai – 400 071
Phones : 022 – 40497401 /02
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Tara Art Printers Pvt Ltd, A-46,47, Sec-5, Noida, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020.

SUBSCRIPTION RATES

Year	Newstand Price (₹)	You Pay (₹)	Overseas
Five	6000	3600	—
Three	3600	2520	—
One	1200	960	US$ 120

Kindly add ₹ 50/- for outside Delhi cheques.
Please send payments only in favour of **EFY Enterprises Pvt Ltd**.
Non-receipt of copies may be reported to **support@efy.in**—do mention your subscription number.

www.indiaelectronicsweek.com

Exhibition & Knowledge Partner
ELCOMA
An EFY GROUP EVENT

A unique platform to **showcase** latest products and technologies in the world of **LEDs & Lighting.**

Only LED show in South India—a market that is known to value technical features and quality.

LEDASIA.in

January 11th-13th, 2016 • Bengaluru

www.ledasia.in

"Make in India" Offer

Special package for those who manufacture LEDs and LED Lighting products in India!

To increase your sales by exhibiting in this event, Call: **Jyotsna (011-40596605)** or write to us at: **efyexpo@efy.in**

Co-located at

INDIA ELECTRONICS WEEK

January 11th-13th, 2016
Bengaluru

NEW PRODUCTS

On-ear headphones from Sony

Sony has recently released a new on-ear headphone called the H.ear MDR-100AAP headset. The device comes with excellent sound quality along with good looks.

The headset is powered by 40mm drivers capable of reaching frequencies of up to 60kHz. It combines premium high resolution audio quality with long hours of listening comfort. The high resolution feature of the headset enables users to hear every drumbeat and note of the music. The device is equipped with lightweight CCAW voice cables and titanium coated diaphragms to minimise unwanted vibrations and enhance sound quality.

Apart from its ergonomic design, the headphones are of a compact folding design that makes them easy to fit in a bag or a coat pocket.

The smartphone compatible H.ear MDR-100AAP comes with a carrying case, in-line remote and a tangle-free cable. The product is available in vibrant colours – veridian blue, red, charcoal black, lime yellow and Bordeaux pink, via online and retail stores.

Price: ₹ 12,990

Address: Sony India Pvt Ltd, A-31, Mathura Road, Mohan Co-Operative Industrial Estate, New Delhi 110044; **Ph:** +(91)-11-66006600; **Website:** http://www.sonyindia.co.in/

Shatterproof display smartphone from Motorola

Price: ₹ 25,000 approximately

After the Moto X Play and Moto X Style, multinational telecommunications company Motorola has launched the Moto X Force, which has been named Droid Turbo 2 in the US market.

The device sports a 13.7cm (5.4 inch) AMOLED QHD display that measures 7.6mm and comes with shattershield technology, the USP of the device. The display is made of five layers of a glass-free composite material that safeguards it from the usual knocks and bumps. It also has a water-resistant nano-coating.

It is powered with a Snapdragon 810 octa-core processor, 3GB RAM and an Adreno 430 GPU. It runs on Android 5.1.1 Lollipop and most probably will get a Marshmellow update soon. The smartphone has a large non-removable 3,760mAh battery, and also supports both Q1 and PMA wireless charging.

On the camera front, the device comes with a 21MP rear camera with f|2.0 aperture, auto-focus, dual LED flash, HDR and 4K video recording along with a 5MP front-facing camera with a f|2.0 aperture.

The Moto X Force has a front-facing loudspeaker along with five microphones for advanced noise cancellation and comes with NFC support. On the connectivity front, the device is enabled with Wi-Fi 802.11, Bluetooth v4.1 and GPS.

Address: Motorola Solutions India, 415/2, Mehrauli-Gurgaon Road, Sector 14, Near Maharana Pratap Chowk, Gurgaon, Haryana 122001; **Ph:** 0124-4192000; **Website:** www.motorola.in/

Dual boot laptop from Datamini

Price: ₹ 9,999

Desktop computer manufacturer Datamini has recently released a two-in-one dual-boot laptop that runs on both Android 5.1 and Windows 10. It features a 25.65cm (10.1 inch) display with a screen resolution of 1280 x 800 pixels. It is powered by a 1.33GHz quad-core Intel Atom Z3735F processor.

The laptop features 2GB RAM, 32GB of inbuilt storage and supports a microSD card up to 32GB. The dual-boot laptop is easy to carry as it comes with a detachable keyboard. It is also equipped with a 5 megapixel rear and 2 megapixel front camera.

The laptop supports the 3G network via SIM slot, as well as Wi-Fi 802.11 b/g/n, Bluetooth 4.0, microUSB and mini HDMI. The device is backed by a 6600mAh Li-ion polymer battery.

The Datamini two-in-one laptop is exclusively available via Snapdeal.

Address: Datamini Technologies (India) Limited, Unit No. 101/104, SDF IV, SEEPZ, Andheri (East), Mumbai – 400096; **Ph:** 022-40428888; **Website:** www.datamini.co.in

An affordable 4G smartphone from Intex

Price: ₹ 10,390

India's second largest smartphone maker, Intex, has launched its latest 4G enabled smartphone – the Aqua Super. The phone comes with a 12.7cm (5 inch) HD IPS display and sports a light body, sleek design and premium black cover.

The 4G capable device is powered by a 1GHz quad-core MediaTek MT6735P processor, runs on Android Lollipop v5.1 and has 3GB RAM. The Aqua Super comes with 16GB inbuilt storage, which can be expanded up to 128GB via microSD card. The smartphone is powered by a 2150mAh Li-ion battery and comes with inbuilt features like Smartwake and Air Shuffle.

The device sports an 8 megapixel rear camera with auto focus, LED flash and a 2 megapixel front camera. The camera features of the phone include face detection, auto scene detection, self-timer, etc.

The Intex Aqua Super is available in black, champagne, grey and white, with a one-year screen breakage warranty via online and retail stores.

Address: Intex Technologies (I) Ltd, D-18/2, Okhla Industrial Area, Phase-II, New Delhi-110020; **Ph:** 91-11-41610224/25/26; **Website:** www.intex.in

Tiny Bluetooth speaker from iFrogz

Price: ₹ 1,631

Smartphone and computer accessories manufacturer, iFrogz, has recently launched a tiny Bluetooth speaker – the iFrogz Tadpole Active.

The device is aimed at providing good quality audio in a tiny package. The Tadpole Active is a lightweight device of just 68gm made of hard plastic on the front and back, with a rubber bezel that runs around it. The dust and water-resistant device has a single button on the left side along with an LED notification light. The multipurpose button can be used to control music as well as calls. A single press will play/pause music, a double press will skip to the next track and a triple press will revert to the previous track.

The notification LED blinks in red when the device is on and turns blue when connected.

The portable Bluetooth speakers can be hooked into the user's belt or keychain through a carabineer (metal clip).

The device comes with a built-in rechargeable 400mAh battery that works for around three hours. It also supports microUSB cables and can run the speaker directly from a power socket or battery pack. The Tadpole Active is available in white, black, blue, green, purple and red via Amazon.

Address: Amazon India, 26/1, 8th Floor, Brigade Gateway, Dr Rajkumar Road, Malleswaram West, Bengaluru, Karnataka – 560055; **Ph:** 91-80-33273000

4G LTE enabled smartphone from Micromax

Micromax has launched a new 4G LTE enabled smartphone called the Canvas Express 4G.

In collaboration with Airtel, the company has bundled a 'Doubled Data' offer to 3G and 4G users. The smartphone has a download speed of up to 150Mbps with Cat.4LTE. It comes with 12.7cm (720 x 1280 pixels) HD IPS display, powered by a 1GHz quad-core MediaTek MT6735P processor along with 2GB RAM. The device runs on Android 5.1 Lollipop, out of the box.

The dual SIM smartphone has 8GB of inbuilt storage expandable via microSD card up to 32GB along with a 2000mAh battery. The device is equipped with an 8 megapixel rear auto-focus camera with LED flash and a secondary 2 megapixel front-facing fixed-focus camera.

Apart from 4G, the device offers 3G, GPRS/Edge, Wi-Fi 802.11 b/g/n, microUSB and Bluetooth options.

The Micromax Canvas Express 4G is available via Flipkart.

Address: Micromax Informatics Limited, 90B, Sector-18, Gurgaon 122015; **Ph:** +91-124-4811000; **Email:** info@micromaxinfo.com

Price: ₹ 6,599

YOU SAID IT

Requesting a digital edition of *OSFY*

As the world is slowly moving towards digital news and books, it would be great if we could subscribe to a digital edition of *OSFY*, so that we can carry it easily on devices like our smartphones or e-readers like Amazon Kindle.

—Venkatapathy Mahalingam, mvenkatapathy@gmail.com

ED: *Thanks for writing to us. We do have the digital edition of the magazine. You can find out more details at* http://efy.in/digitalpublications/. *I would also like to mention that if you subscribe for the print edition, you get the login details for the e-zine edition too. You can buy or read the OSFY e-zine at http://ezine.lfymag.com/buy/purchaseezinesingleissue.asp*

Suggestion for using CentOS in place of RHEL

In the October 2015 issue of *Open Source For You,* the admin section has an article on Chef ('Chef: Recipes that turn infrastructure into code'). The author asks us to install RHEL, which requires a paid subscription. Instead, we can install CentOS in its place since it's completely free. I thought it may be helpful to point this out in the coming articles. Hope this helps.

—Gautham, angautham@yahoo.com

Hi Gautham!

Though we have explained with RHEL as an example, the steps are pretty much the same in CentOS as well. This is a series on Chef and we will be referring to CentOS in forthcoming articles. Thanks for the feedback.

Author: Magesh, magesh.kasthuri@wipro.com

ED: *We thank you for your suggestion, which has been accepted by the author as well. Hope to see this implemented in future articles.*

Answers to CodeSport questions

I subscribe to and have been reading *OSFY* for more than two-and-a-half years now. It is a great magazine to learn about the latest developments in the open source world. I like the CodeSport section by Sandhya Mannarswamy very much and try my luck with all the questions.

It would be really helpful if the answers, links or best solutions are published in the subsequent issue of the magazine. Currently, there is no way we can validate our answers. I really look forward to knowing the best solution for a given problem.

—Chethan, ChethanKumar.C@LntTechservices.com

ED: *Thanks a lot for the feedback. It really makes us feel good to know that you like the articles published in the magazine. We value our readers' suggestions as they help us to improve the quality of the magazine. It feels great when readers tell us how they benefited from our magazine. We will convey your feedback to the author of the article and will surely try and implement your suggestion.* END

Please send your comments or suggestions to:

The Editor,
Open Source For You,
D-87/1, Okhla Industrial Area, Phase I,
New Delhi 110020, **Phone:** 011-26810601/02/03,
Fax: 011-26817563, **Email:** osfyedit@efy.in

ESDS bags cloud computing patent from USPTO

Exuberant Support for Data Services (ESDS), a leading cloud services provider offering managed data centre services, has received India's first patent in this domain for its eNlight Cloud from the United States Patent and Trademark Office (USPTO).

ESDS is recognised for its innovation that promises to completely transform commodity based computing into a utility model. This is achieved through auto-scaling of computing resources that deliver 100 per cent availability of customer applications on a pay-per-consume model.

According to US based research firm, Gartner, the present average global compute consumption by any organisation is extremely low at 8 per cent while it is just 5 per cent among Indian organisations. Unlike the currently prevalent pay-per-use cloud model, ESDS' eNlight Cloud delivers inbuilt intelligence to constantly monitor consumption of virtual machines and, accordingly, assign computing resources in real-time.

Organisations ranging from large enterprises to SMEs are set to significantly benefit from the on-demand availability for maximum computing loads, enabling them to make IT investments based on their actual requirements and not on their financial strength.

"With our strong legacy in managing thousands of virtual machines, ESDS is committed to delivering innovative cloud technology that is scalable and affordable. Additionally, this patent is a testimony of our commitment to deliver on the government's 'Digital India' vision by making IT compute accessible and affordable to all businesses," said Piyush Somani, MD and CEO of ESDS.

Citrix forms OpenStack partnership with Red Hat

Citrix is introducing the integration and certification of its application delivery controller, NetScaler, with the Red Hat Enterprise Linux (RHEL) OpenStack platform. "For the first time, customers can assemble their cloud infrastructure using best-of-breed components from Citrix and Red Hat," sources from the companies claim.

The partnership aims to provide clients with access to a broader portfolio of products to further boost the performance and scale provided by RHEL OpenStack platform deployments.

"Customers choose OpenStack because they want a truly open, production-ready platform for building their public or private clouds," said Radhesh Balakrishnan, GM of OpenStack at Red Hat.

"NetScaler is now certified to work with the RHEL OpenStack platform, and offers customers the rich capabilities of NetScaler for workload availability, security, scalability, and performance," he added.

Arch Linux 2015.10.01 ISO supports Linux kernel 4.2

The major new modification in the Arch Linux 2015.10.01 ISO image is its

Mozilla offers fund to open source projects

Mozilla CEO Mitchell Baker has announced 'a new level of support' in the form of US$ 1,000,000 to be distributed in the form of awards among 10 yet to be determined open source projects by December 12.

"The Mozilla Open Source Support programme is intended to identify and commemorate communities who are leading the way with open source projects that contribute to our work and the health of the Web," said Baker.

It encompasses two elements:

(a) A 'give back' element for open source and free software projects that Mozilla relies on.

(b) A 'give forward' component for supporting other projects where financial resources from Mozilla can make the entire community more successful.

The initial 10 projects will be part of the 'give back' part, added Baker.

The specific terms of the programme are still being formulated and members of the open source community are invited to offer their suggestions.

"I am reminded repeatedly of how deeply Mozillians identify open source and free software as a significant element of an open Internet and healthy, trustworthy online experiences," said Baker. "I am excited to build a programme that helps us bring concrete support to this worldview," he said.

NIIT launches next generation programming

NIIT Limited has launched what it claims is next generation programming for developers. The 'Programming using Python' course has been structured such that even a student with no programming knowledge can easily grasp the language. This course is also ideal for students who have worked with Java or C# and now wish to enhance their skills by learning another popular open source language.

"This course will be transformational for students who want to become successful Python programmers—it aims to create skilled programmers for the industry," said Shivan Bhargava, group president-skills and careers group, NIIT Ltd. It will help the students in constructing program blocks using functions, and it will also help in understanding the construction of networked programs using Web services.

Python is considered as an engaging programming language to learn, because of its interesting syntax features. It is considered as one of the most preferred options today for any high performance computing.

The syntax of the language is designed to be readable. A recent study has shown that Python is one of the most in-demand programming languages and is ranked second in the 'Top 10 Popular Programming Languages of 2015'.

support for the Linux 4.2 kernel, as the September build was driven by Linux kernel 4.1.6 LTS. Users can now set up Arch Linux with Linux kernel 4.2.2. Moreover, the systemd init system was restructured to version 226.

One can get the new ISO image, Arch Linux 2015.10.01 via Softpedia or straight from the distribution's website, where the GNU/Linux operating system is offered for download as a dual-arch ISO-hybrid image that ropes in both 64-bit (x86_64) and 32-bit (i386) hardware architectures.

Existing Arch Linux users don't need to download the new ISO image to keep their systems efficient. It is a lightweight and independent operating system that follows a rolling release model. In order to have all the latest updates that are being launched, one should run the 'sudo pacman –Syu' command in a terminal emulator application.

The Arch Linux 2015.10.01 ISO image is only for those who want to either reinstall their Arch Linux operating system from scratch or set up the GNU/Linux division on new computers.

LFA now based on Ubuntu 15.10, Debian 8.1 and Fluxbox

The Linux For ALL (LFA) Build 151024 has been entirely restructured and it is now based on the recently launched Ubuntu 15.10 Wily Werewolf OS and the Fluxbox window manager. However, it still includes various essentials from the Debian GNU/Linux 8.1 (Jessie) distribution.

The LFA Build 151024 updates all the major installed components and applications to their latest versions available as on October 24, 2015. It is powered by Arne Exton's special 4.2.0-16-exton kernel, which is based on the upstream Linux kernel 4.2.3.

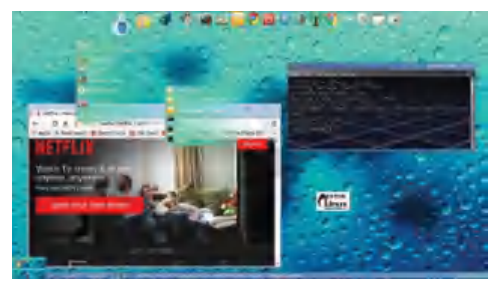

"The latest upgraded version of LFA is now available and all the installed packages have been restructured to the latest edition," Arne Exton said in an email to Softpedia.

The version consists of Google Chrome, SMPlayer and Synaptic Package Manager.

The software collection of the Linux For ALL distribution has always been grand and the most recent version comes with the Google Chrome Web browser, which can be used to watch Netflix movies, apart from the Synaptic Package Manager for installing extra software, and the powerful, MPlayer-based SMPlayer video player.

The distribution also includes the Cairo-Dock dock-like application launcher, the Dreamweaver-like BlueGriffon HTML editor, a GParted partition editor, a PCManFM file manager, an Adobe Flash Player plugin and all the other tools required for assembling programs from sources.

Developers working on Raspberry Pi 2 sharing platform

The developers of Ubuntu MATE are working on some specific tools that will allow them to build the distribution for the Raspberry Pi 2 platform. The researchers are looking to distribute these tools so that other projects can use them, like Xubuntu or Lubuntu. Creating the distribution for Raspberry Pi 2 might become a routine thing now after the Ubuntu MATE team has put it on its release schedule.

The interesting thing is that the Raspberry Pi 2 image is not part of the customary release, but that didn't stop the team at Ubuntu from doing it anyway.

The developers have collectively put together some tools and now they want a little aid from the community to name them.

There are a number of Linux distributions on hand for Raspberry Pi 1 and 2, but the community always wants more. There are some very good technical reasons why we won't see Ubuntu or Kubuntu running on a Raspberry Pi shortly, but those reasons don't apply to all the official flavours.

The developers are setting an example in this regard, and it looks like they already managed to release one of the best distros for this mini PC. The tools are already under production, and anyone interested can download the source package for it, but please keep in mind that they still need to be enhanced. In any case, if things go well, we should see a series of new distros for Raspberry Pi 2 soon.

In the meantime, you can download the Ubuntu MATE 15.10 for the Raspberry Pi distribution from the authorised servers.

Netflix updates its open source policies

Video-streamer and entertainment production house, Netflix has restructured its open source policies. One remarkable change is a decision to release code pre-packaged in Docker's container formats.

Netflix has open sourced a lot of its production code, but now says those efforts were puzzling some because it was getting harder to figure out which projects were useful for a particular company or a team. "The external community was also uncertain about which components Netflix continued to put in and support, and which were in maintenance or sunset mode," developers reported.

A revival of the company's GitHub outpost is among its other developer-friendly efforts, which include a pledge to do the four following things from now on:

- It will provide full transparency on which projects are archived, i.e., no longer actively developed or maintained. It will not be removing any code from GitHub repos, but will articulate if the company is no longer actively developing or using a particular project. This will affect and reflect on its OSS projects as Netflix needs to change over time.
- Offer an enhanced roadmap for what new projects the company is planning to open, and which open projects are still in a state of heavy change. This will allow the community to better decide whether particular projects are appealing.
- Expose some of the internal metrics Netflix uses to evaluate its OSS projects - number of issues, commits, etc. This will provide better transparency regarding the maturity/velocity of each project.
- Docker is gradually becoming more accepted and Netflix is a prominent company, so the decision to Dockerise could well create ripples and additionally, move forward the container platform's popularity and reach – if at all that's possible, given that Microsoft, VMware, OpenStack and many others have all decided that Docker's an impressive way to do containers.

Neo4j unveils graph query language called openCypher

The Neo4j graph NoSQL database team has unveiled an open source graph query language called openCypher. Neo Technology announced the launch of the open source project, which will be accessible to technology providers as a common language for querying graph data. It is based on Neo4j's query language that is used to accumulate and recover data in the graph database.

At present, there is no common query language standard for working with graph databases, like SQL for accessing data in relational databases.

The objective of openCypher is to speed up the growth of graph processing and analysis. Technology providers can implement openCypher inside their tools and platforms. It delivers four major artifacts under a permissive licence:

Language specification: The Cypher language specification is a technical expression of the language syntax to enable parsers to auto-generate the query syntax. A full semantic specification is also planned as a part of the openCypher project.

Reference implementation: Distributed under the Apache 2.0 License, reference implementation is a fully functional implementation of key parts of the stack needed to support Cypher inside a data platform or tool. The first planned deliverable is a parser that will take a Cypher statement and parse it into an abstract syntax tree representation.

Technology compatibility kit (TCK): This consists of tests that software vendors can run on their products, to self-certify support for a given version of Cypher.

Cypher reference documentation: This includes the user documentation describing the use of the Cypher query language with examples and tutorials.

A number of organisations are supporting the openCypher project and these include Oracle, Databricks, Tableau and Linkurious.

You can learn more about openCypher by going on the website's FAQ page.

HP creates open source NOS

HP has joined hands with an array of firms, including Arista, Broadcom and Intel, to produce OpenSwitch, an open source NOS (network operating system). Developers can now pool resources, test new theories and innovate to build up higher-quality networks for companies. Moreover, these networks can be tailored to accommodate specific business needs.

The digital character of the world has produced a need for more compliant NOSs. Networks handle hefty amounts of data every day, which has led to a need for on-demand scalability.

Conventional networks are built with proprietary software, which doesn't permit customers or their software partners to transform their networks to their own parameters.

But now developers can focus on improving their business-specific workload needs and functions instead of spending time and energy struggling with the complex licensing structures.

The primary benefit of OpenSwitch is that it allows developers to create new applications and specific features for individual networks, and it also helps them by enabling these resources to finish the project faster and with fewer post-release tribulations.

Open source projects attract a huge number of developers' contributions from around the globe, and hence such projects are bound to be less tricky and get completed faster.

HP Helion OpenStack 2.0 offers a cloud platform

HP has launched the production ready HP Helion OpenStack 2.0, an open source based cloud platform designed to meet enterprise requirements. Many companies are turning to the flexibility and economics of the OpenStack project, as they struggle to find the right mix of conventional IT and private cloud technology to run their mission-critical applications and protect sensitive data. The HP Helion OpenStack 2.0 offers an enterprise grade cloud platform, while adding new features to address companies' lifecycle management and security challenges. It provides trouble-free provisioning of new infrastructure and the capacity to repurpose existing infrastructure to assemble scalability requirements without impacting availability.

The rolling upgrades smoothen the process of the entire cloud environment software upgrades without the need for planned or unplanned downtime. And its continuous patch management allows security patches and updates without application interruption.

HP Helion OpenStack 2.0 also enables customers to create and manage software defined networks (SDN) in a distributed, multi-data centre environment through integration with HP Distributed Cloud Networking (DCN) and Nuage Networks Virtualized Services Platform. In addition, it removes the boundaries of conventional networking, unlocking the full mechanisation and agility needed for the hybrid cloud.

"Enterprises want to benefit from the powerful capabilities of OpenStack technology, but they must have the enterprise grade capabilities required to support their businesses," said Narayanan Chellappan, country director, cloud division, HP India.

HP Helion Professional Services can offer a team of qualified HP architects and cloud technologists to help clients decide on the right cloud strategy. This brings an exceptional perspective to the customer, gained from delivering thousands of lines of code to the OpenStack community, as well as hands-on customer deployments of large scale OpenStack clouds. These services also have capabilities that span the entire cloud lifecycle, from advice and strategy, to transformation, implementation and management.

Red Hat delivers latest, stable developer tools

Red Hat Inc. has announced the common availability of Red Hat Software Collections 2.1, the latest, stable open source Web development tools, dynamic languages and databases. These enhance developer agility by helping to speed up the creation of modern applications that can then be more confidently deployed.

Red Hat Developer Toolset 4 is also by and large available, which helps to make application development more efficient on Red Hat Enterprise Linux (RHEL) by giving developers access to the latest open source C and C++ compilers, and complementary development and performance profiling tools.

Red Hat Developer Toolset enables developers to assemble applications once and deploy them across several versions of RHEL. It is accessible through the RHEL Developer Program and associated subscriptions.

Red Hat Software Collections 2.1 features a number of new and restructured collections. These are:

1. Varnish Cache 4.0.3, a caching HTTP reverse proxy that speeds up Web applications
2. nginx 1.8, a simplified version of the trendy HTTP and reverse proxy server

3. node.js 0.10.40, a restructured version of the runtime environment for developing server-side Web applications
4. Simplified Maven 3.0 compilation with additional packages

New additions and simplified components of Red Hat Developer Toolset 4 include:

1. GNU Compiler Collection (GCC) 5.2
2. The GNU Project Debugger (GDB) 7.10
3. Eclipse Mars (4.5)
4. Several updates to toolchain components and performance tools, such as binutils (2.25) and SystemTap (2.8)

A lot of the popular Red Hat Software Collections have been made available as Docker files for developers looking to gain from the rapid progress and deployment cycles inherent to Linux containers. Many of the popular Red Hat Software Collections are also now on hand as Docker-formatted container images via the Red Hat Customer Portal. Applications built with Red Hat Software Collections can be deployed in production with greater assurance because most software collections and components are supported for three years.

Netflix launches open source Spinnaker for AWS

Netflix has launched its open source Spinnaker project after more than a year. It is designed as a substitute to the Amazon Web Services (AWS) cloud management system, Asgard.

It has been developed along with Google, Microsoft and Pivotal. The 'continuous delivery platform' lets companies hook into and arrange resources across two cloud providers, simultaneously.

"Spinnaker facilitates the design of pipelines that represent a delivery process that can begin with the creation of some deployable asset and end with a deployment," Netflix's manager of delivery engineering, Andy Glover, said in a company blog post.

According to the company, Spinnaker offers an "…out-of-the-box setup," and according to Google, engineers can make and reuse the pipelines on different workflows.

Meanwhile, Netflix is working on Microsoft Azure support, which will follow Google Cloud Platform, Cloud Foundry and AWS. For Asgard users, the existing assets are compatible with Spinnaker.

Visual Studio now supports debugging of Linux apps

Global tech giant Microsoft has released a preview of a Visual Studio extension that includes remote debugging using GDB (GNU Project Debugger).

Developers can now debug applications running on Linux servers or IoT devices from the comfort of Visual Studio. The company aims to give its developers the broadest platform it's ever had, which will be capable of handling Android, iOS and Linux development, alongside the expected Azure, Office and Windows.

Microsoft is aiming to make Visual Studio the finest development environment around. The free and cross-platform Chromium-based code editor Visual Studio Code has been open sourced. A new build is also available, adding an extension mechanism to the editor. As of now, about 60 extensions are on hand, together with new language support for Go, richer debugging, code linters, etc.

Extensions for Visual Studio Code are found in a new place for distributing both free and paid extensions: Visual Studio Marketplace. This will substitute the old Visual Studio Gallery and will be a single place for finding extensions for all versions of Visual Studio (community, professional and enterprise),

Raspberry Pi may soon become more Ubuntu-friendly

Raspberry Pi is likely to become more Ubuntu-friendly as new tools make it easier to port certain versions of Canonical's operating system to it. These tools were created by the developers of Ubuntu MATE, which uses the MATE desktop environment as a substitute for Unity.

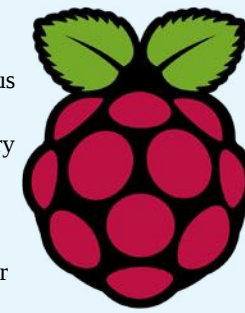

At present, Raspberry Pi devices are minus the graphics support necessary to run the most complicated desktop environments for Ubuntu, such as Unity, GNOME and KDE.

Developers have already made a range of Ubuntu flavours, including server versions, for the Raspberry Pi 1 and 2. The Ubuntu flavours are intended for those interfaces that won't work on the Raspberry Pi, except if users are content working only from the command line. However, lighter-duty desktop environments run in complete graphics mode on the devices, which means that Lubuntu and Xubuntu, which make use of the LXDE and Xfce desktop environments, work well on the gadget when compiled using the MATE team's porting tools.

Ubuntu Server also works on the hardware, which doesn't include any graphical interface, by default. This development is expected to lure only Ubuntu power users and programmers who would like to build a more comfortable ecosystem of Ubuntu software for the Raspberry Pi.

Meanwhile, the porting tools may show the way to bigger developments down the line, predominantly in the sphere of Ubuntu applications and programs for rooted or modest devices.

Visual Studio Team Services (formerly known as Visual Studio Online) and Visual Studio Code.

The global tech giant will also be using Marketplace to sell a new kind of Visual Studio subscription. Visual Studio Professional and Enterprise are together obtainable through MSDN subscriptions bought on an annual basis. These subscriptions also provide development access to many other pieces of Microsoft software.

The company is now offering Visual Studio Professional and Enterprise as cloud subscriptions, with the option of buying these either each month or annually.

The company will unveil a new Visual Studio bundle called Dev Essentials, to make all-device development easier to access like Visual Studio Community Edition, the free tier of Visual Studio Team Services, and from early next year, monthly Azure credits. It will supposedly comprise everything needed to develop for Windows, Android and iOS.

Visual Studio

However, building for iOS still requires access to a Mac; Visual Studio remotely controls the OS X toolchain to perform the actual building and application deployment. Microsoft has partnered with MacinCloud to provide access to OS X build agents for US$ 30/month for those who don't want to manage Macs on their own.

In addition to using the built-in support for Android and iOS development, Visual Studio is often used in combination with the Xamarin platform for cross-platform applications. The release of Xamarin 4 comes with a new crash and performance analytics tool, Xamarin Insights, and version 2 of its Xamarin Forms library. It adds new iOS 9 and Android Marshmallow controls, and a new build agent for performing software builds on OS X from inside Visual Studio that should be easier to set up and more reliable than the current mechanism.

As well as the GDB debugging extension, Azure Service Fabric—the management infrastructure for microservices on Azure that is used by Microsoft for Cortana, Lync, and Azure SQL, is now accessible for public preview.

Canonical releases OpenStack Autopilot

Canonical has released OpenStack Autopilot, a cloud management tool for Ubuntu Linux in a bid to bring OpenStack private clouds to the masses. This is Canonical's latest move in the open source cloud computing market.

Till last year, Autopilot was available in beta form, but it is now generally available and is a feature in Landscape, Canonical's platform for managing deployments of Ubuntu systems.

It extends the product to include support for automated OpenStack deployment and administration using Ubuntu servers. It can set up clouds, add new hardware to existing clouds and assist with cloud management.

Autopilot also offers the ability to run OpenStack administrative services inside dedicated containers via LXD, Canonical's home-grown virtualisation hypervisor.

According to Canonical, it supports a range of software-defined storage and networking options which now include OpenDaylight, an open source SDN that is new to Autopilot, debuting with the general-availability release.

The company is offering Autopilot for entities to build private clouds using OpenStack, without having to spend money on hiring OpenStack knowhow.

CODE SPORT

Sandya Mannarswamy

In this month's column, we feature a set of computer science interview questions.

Over the last couple of months, we have been discussing a set of questions on machine learning and data management. This month, we return to our discussion of topics in natural language processing (NLP). We focus on the important problem in NLP known as parts of speech (POS) tagging, covering why it is important and some of the techniques used in this area.

Parts of speech tags are used to tag words in a sentence with their syntactic category, such as whether a word is a noun, verb, adjective or adverb. For instance, given the sentence, 'The dog barks,' its POS tags would be as follows: 'The (determiner) dog (noun) barks (verb).' The tags in brackets indicate the parts of speech associated with the word in the sentence. Some readers may wonder why one needs to do POS tagging. Isn't the POS tag of a word known automatically? Well, the answer is not so straightforward.

Most of the words in the English language (and in the case of most other languages as well) can take different tags based on the context in which they are used. For example, consider the two sentences:

(1) Please give me the book.

(2) Please book the evening flight to Delhi.

In Sentence 1, the word 'book' is associated with the noun POS tag. In Sentence 2, the same word is associated with the verb POS tag. Certain tags such as prepositions in English are easier to tag in a sentence, since they are associated with an unambiguous POS tag. On the other hand, most other words have an ambiguous association with POS tags, since they can take different tags based on the context. POS tags can be divided into two broad categories, namely, Closed Class Tags and Open Class Tags. In the first category, member words are typically fixed and form a closed set. Open Class Tags are those whose membership can increase over time due to the addition of new words. For example, a preposition is a Closed Class Tag since, typically, in English, new prepositional words don't get added. On the other hand, consider tag classes such as nouns and verbs. In English, new nouns can get added due to new entities such as 'Google' or 'iPhone' and new verbs can get added, such as 'fax', etc. Hence, verbs and nouns are Open Class Tags.

Apart from POS tag classes such as nouns, verbs, prepositions, etc, how many POS tag classes are typically considered for POS tagging of English sentences? While there are different tag sets available with differing number of tags, a popular one in the NLP community is the Penn Tree Bank tag set of 45 tags, available at: *https://www.ling.upenn.edu/courses/Fall_2003/ling001/penn_treebank_pos.html*

To those still wondering why POS tagging is needed for text, it is a pre-processing step that's necessary in a number of related problems, namely, information retrieval (IR), text-to-speech processing, word sense disambiguation, etc. POS tagging is also useful for syntactic parsing, since the tag associated with a word can be used to infer syntactic classes of surrounding words. For instance, nouns are preceded by determiners and adjectives, and verbs are typically followed by nouns, and so on. There has been considerable prior work in using POS tagging to improve IR performance, by cutting down the terms indexed based on their POS tags. For instance, since indexing only the nouns in a document generally provides a good idea on what the document is about, it is possible to create a pruned index set using POS tags. Similarly, POS tags can be used to improve text-to-speech systems, in deciding how a word should be spoken with different intonations. For instance, consider the following two sentences:

(1) I object to your insolent tone.

(2) What is the object of your shipping expedition?

In Sentence 1, 'object' being tagged with the verb POS tag would be spoken with a different intonation as opposed to its usage in Sentence 2, where it is being used as a noun in the query.

POS tagging is also useful in the important but challenging problem of word sense disambiguation. Polysemy of many languages means that a single word can have different meanings associated with it. POS tagging of the sentence can aid in word sense disambiguation. Consider the following two sentences:

(1) I bank with HSBC.

(2) I sat on the river bank.

We can use the POS tags to disambiguate between the different meanings associated with the word 'bank'.

Given the usefulness of POS tagging, now let us discuss the possible approaches to it. First, let us define the problem of POS tagging formally. Given a sentence consisting of words {w1, w2, w3,… wn} and a tag set T = { t1, t2, … tm}, the task is to identify the best possible tag sequence that can be assigned for the sentence S.

Just as in the case of various other tasks in NLP, POS tagging can also be performed using both supervised and unsupervised learning techniques. Supervised POS tagging requires that you are given a text Corpus C and a Tag set T. The Corpus C consists of a set of text documents and each document is a collection of sentences. Each Sentence S is a collection of words. Each word in S is tagged/labelled with the appropriate tag set element from T, resulting in a tag sequence <t1,t2… tn> being associated with the Sentence S consisting of <w1, w2, … wn>. Hence, we are given a set of <sentence, tag sequence corresponding to the sentence> as training data. The task is to learn a function that can predict/assign the tag sequence corresponding to the input test sentence, given the labelled training data.

While supervised POS tagging approaches perform well in terms of accuracy, the disadvantage is that they need the corpus to be labelled with the tags. Unsupervised POS tagging is based on finding equivalent word classes; however, without any labelled data, it cannot assign human readable tags to words. Hence, a small amount of manually labelled corpuses can be used with unsupervised POS taggers to induce a mapping between human readable tags and the word equivalence classes induced by the tagger. We will discuss the unsupervised POS tagging in detail in the next column.

Another way of categorising the POS tagging approaches is as either rule based POS tagging techniques or stochastic POS tagging techniques. Rules can either be provided manually by linguistic experts, or they can be automatically induced by learning from large text corpuses which are tagged. The rules can then be used to tag an untagged sentence. Stochastic POS tagging techniques are based on finding the probabilities of the tag sequences associated with a given sentence. Assume that you have a black box, which can assign a probability of 'p[i]' to each possible tag sequence 'i' for sentence S. Given these probabilities, the stochastic approach can choose that tag sequence whose 'p[i]' is the maximum from among all possible tag sequences.

Now let us consider that we have a sentence S whose word sequence is <w1, w2, .. wN>. Each word 'w' can be tagged with one of the possible tags from tag set T. Let us assume that tag set T contains M tags. Hence, each word can be assigned one of M possible tags. Can you calculate how many tag sequences are possible for the given sentence S?

The number of possible tag sequences associated with S becomes very large and is equal to MN. Hence, brute force search among all possible tag sequences is not efficient. Even if the probabilities are known for each possible tag sequence associated with S, it is computationally expensive to find that particular tag sequence whose probability is the highest.

Now can you think of ways by which we can avoid this brute force search? The question to our readers is this: Can you think of any solution from your algorithms course, that you have used in similar problems? For instance, consider the example of finding the shortest path from a source vertex to a destination vertex in a graph. Since there are exponentially many possible paths, a brute force listing of each path and its associated cost is computationally expensive. In such cases, can you recall what kind of algorithmic techniques you used to overcome this issue? A clue I would like to give you is to think about the dynamic programming approach. We will continue the discussion of POS tagging approaches in our next column.

If you have any favourite programming questions/software topics that you would like to discuss on this forum, please send them to me, along with your solutions and feedback, at *sandyasm_AT_yahoo_DOT_com*. Till we meet again next month, happy programming! END

By: Sandya Mannarswamy

The author is an expert in systems software and is currently working as a research scientist in the Xerox India Research Centre. Her interests include compilers, programming languages, file systems and natural language processing. If you are preparing for systems software interviews, you may find it useful to visit Sandya's LinkedIn group Computer Science Interview Training India at *http://www.linkedin.com/groups?home=HYPERLINK "http://www.linkedin.com/groups?home=&gid=2339182"&HYPERLINK "http://www.linkedin.com/groups?home=&gid=2339182"gid=2339182*

Going Against the Tide of the Apps

This article by a veteran of the computing world is an interesting account of his experiences with computers, right from punch cards to mobiles and tablets. Since he has reservations about the apps that are constantly being thrown at him from various websites, he puts forward a case for using a browser through a proxy server, rather than installing an app.

As yet another website nags me to install its app on my Android tablet, I recall the various user interfaces I have gone through.

My interaction with computers began with punched cards. Interactive experience started with a teletype, which was soon improved on by a much faster one, at 300 baud and a much quieter electronic teletype. We graduated to dumb terminals. Then, the most dramatic change came with smart terminals. We were now operating with a whole screen instead of a character at a time. Form based interactive computing with the mainframe and minicomputer applications became possible.

Once the PC arrived, the first connection to the host computers was using a terminal emulator application. With graphical displays on PCs and better networking, we were soon deploying client-server applications. The interaction of the user with the computer became much easier. Consistency of common usage patterns across applications made it much easier to use new applications without the need to learn the new applications' interfaces and commands. The biggest challenge with this model, though, was the distribution of new versions with bug fixes and new features.

Web interfaces were a remarkable relief. It was possible to connect to any site without having to load a specific application. For a number of years, the users of open source software had to suffer. Some websites, especially those related to finance which we could not avoid, seemed programmed only for the quirks of Internet Explorer. However, with the popularity of Firefox and Google Chrome, it is now rare to come across a site that gives users a headache.

Smartphones introduced the challenge of smaller displays to websites. And for some of you, this display on the smartphone, may well be the first interface you've ever experienced.

Websites detect the browser being used and send a different page display, based on the browser. The websites can just as easily detect the device and send different versions of the Web page. However, that has not been a popular solution. The convention of *m.site-url* instead of *www.site-url* for sites optimised for mobile screens is common but not universal.

The apps

A more common solution has been that each site tries to offer an optimum solution by creating apps targeted at different mobile computing environments, identical to the client-server era. Some like Firefox OS, Ubuntu Touch or Blackberry will remain neglected. Your experience may differ based on the platform you are using and the version of the application you have installed.

This brings me back to the irritation with the websites nagging me to install apps. If you click on a link, will it open in an app or in a browser? Even if this step is transparent, it makes no sense to have thousands of apps installed for the countless sites you may browse.

In most cases, you may find it pretty hard to notice the advantage a custom application would offer, compared to a reasonably well defined page for the smaller displays. In fact, how would an app deal with an Android tablet with a screen that is over 50.8cm (20 inches) like the Nabitablet?

Using a proxy server

The advantage of a Web browser over an app becomes very obvious if you wish to use a proxy server. If you are using multiple devices, including a laptop, a desktop, a tablet or a smartphone, and have a limitation on your data plan, it would make sense to reuse the pages and objects which have already been downloaded on one device when you access these from another. Aside from saving data, it will lead to a better performance.

You need to 'long press' on a Wi-Fi connection on a

smartphone and under *Advanced options*, configure the proxy. It is very simple, except that a message warns you that the proxy will be used by the browser but may not be used by the apps.

While I was monitoring the Squid logs, everything seemed fine. Using various Google apps went via the proxy. And then I noticed that one of the biggest users of data was missing – Facebook!

For the first time on the tablet, I used the browser to go to the Facebook site, which as expected, advised me to have a better experience by installing the app. To my surprise, I found the experience of Facebook's mobile website to be far easier and snappier for my usage, than the app. Besides, as expected, these calls to the Facebook site were now via the proxy.

The major shock for me was that I had become so attuned to the 'app' environment that it had never even occurred to me to try the browser instead!

Also, the app stores build up proprietary environments, creating very high barriers for new approaches to the mobile platform. The client-server era made Windows the dominant desktop environment and it continues to dominate even today. Now, environments like Firefox OS have trouble even gaining a foothold in this domain, leave alone getting a chance to change the world!

So, I would encourage you all to think before you install the next app. Is it in your best interests or is it in the interest of the site owner? What benefit does it bring you and at what cost? I am not talking about the cost of the app – but the hidden cost of your data being made available to the site owners! END

By: Dr Anil Seth

The author has earned the right to do what interests him. You can find him online at *http://sethanil.com* and *http://sethanil.blogspot.com*, or you can reach him via email at *anil@sethanil.com*.

Creating and Running Apps on the Microsoft Azure Hybrid Cloud

Microsoft Azure is a cloud computing platform developed to build, deploy and manage Web applications through cloud services. It provides both PaaS and IaaS services. Besides Microsoft software, it supports various programming languages, tools, frameworks and other third party systems and software.

Cloud computing is currently one of the key domains of research. It is used for the delivery of assorted computing services including telecommunications, storage, platforms, software, virtual infrastructure and many others. Cloud computing offers the delivery of virtual resources without the need to worry about establishing and maintaining the actual infrastructure. The end users of cloud services access the remote computing resources using Web based clients without having very complex configurations in their own devices. Those with just 1GB RAM in their own devices can work on the 128GB RAM installed in the cloud server without any issue. In this case, the system or device of the user acts as the terminal that accesses and works on the remote cloud infrastructure.

Cloud computing delivery channels

Cloud computing services are delivered by the service providers using different specific models, as listed below.

Infrastructure as a Service (IaaS): The IaaS cloud involves the delivery of computing infrastructure such as a virtual machine disk image library, raw block storage, object storage, firewalls, load balancers, IP addresses, virtual local area networks, etc, on demand, from the service providers' vast pool of computing resources installed in data centres.

Platform as a Service (PaaS): The PaaS models involve cloud service providers delivering a complete computing platform, including operating system, a programming language execution environment, an integrated development environment, database server, Web server and other related technologies. Popular PaaS implementations include Microsoft Azure and Google App Engine.

Software as a Service (SaaS): In SaaS, also known as on-demand software, the cloud users are provided access to application software and databases. Cloud providers manage the infrastructure and platforms that run the applications.

Metal as a Service (MaaS): This is a provisioning construct developed by Canonical, the developers of Ubuntu, to assist and integrate the deployment and dynamic provisioning of hyperscale computing environments such as Big Data workloads and cloud services, in a dynamic manner.

Network as a Service (NaaS): This service enables the end user to create the virtual network on the cloud without physical installation using hubs, switches and other physical resources. Using this approach, any number of computers can be attached and worked on at remote data centres.

Hybrid cloud services

Hybrid cloud services enable people to use different types of cloud services simultaneously, with just a single point

of execution. Microsoft Azure and Red Hat OpenShift are prominent cloud service providers, and allow users to work on the following:

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS)
- Network as a Service (NaaS) and many other models

Microsoft Azure as a hybrid cloud

Microsoft Azure is one the most effective cloud service providers, and delivers a number of computing resources on demand at very competitive prices. It provides the following computing and technology services:

- Web application development
- Mobile application development
- Data and storage virtualisation
- Internet of Things (IoT)
- Remote network infrastructure
- Content delivery
- Active directory services
- Multi-factor authentication
- Business analytics
- NoSQL databases
- API management
- Visual Studio online
- Media intelligence
- On-demand live streaming and many other services

Visual Studio can be accessed from MS Azure as a cloud service to create and deploy dependent applications. In this scenario, the programmer can work online on Visual Studio, rather than obtain a licence for an individual or specific user.

Figure 1: Microsoft Azure portal

Figure 2: Login panel to create an account on MS Azure

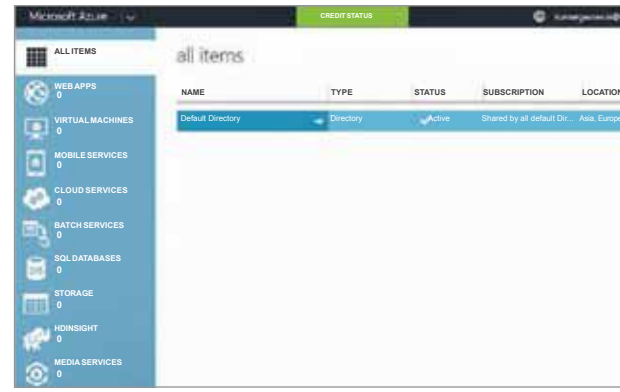

Figure 3: Azure's dashboard to view different services available on demand

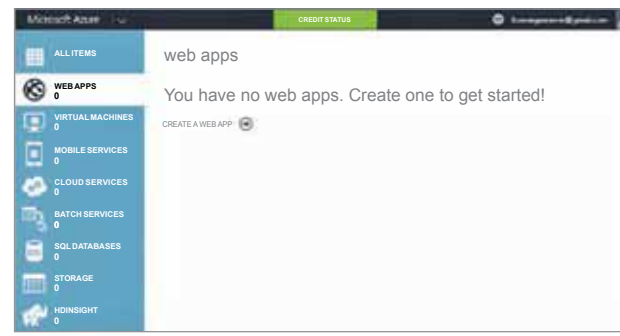

Figure 4: Creating a new Web app in MS Azure

Figure 5: Selecting a plan and region to deploy the Web app

Figure 6: Installing pre-built apps for e-commerce, CMS, blogs and others

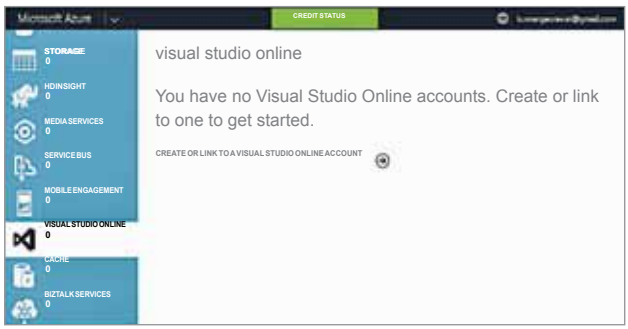

Figure 7: Working on Visual Studio online as Platform as a Service

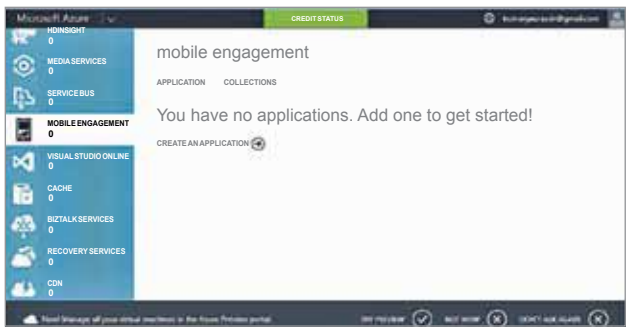

Figure 8: Creating apps on the mobile engagement panel

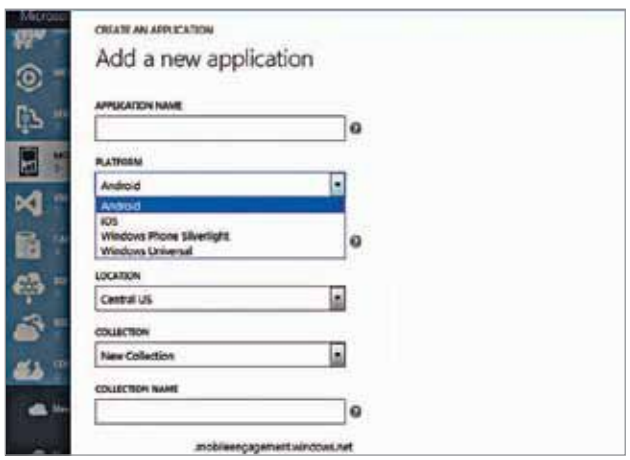

Figure 9: Selecting the type of mobile app on Azure

Figure 10: Creating Web apps using Joomla CMS

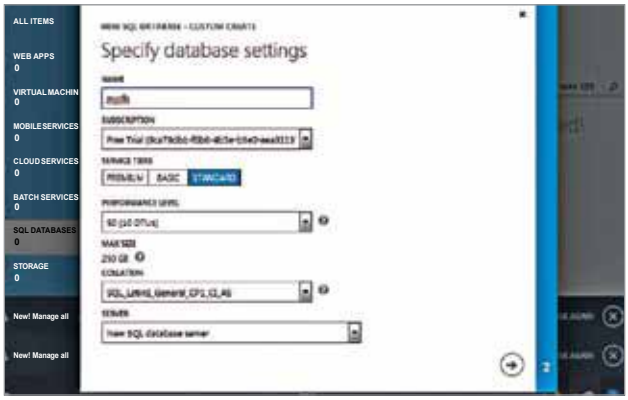

Figure 11: Specifying the database settings for Joomla on Azure

Figure 12: Creating an e-commerce application using OpenCart on MS Azure

As shown in Figure 8, there is the option to create and deploy mobile applications online using the cloud platform. This service is known as PaaS, as MS Azure is providing the complete platform to develop mobile applications on the cloud.

Azure can be used to create and deploy Web applications that are fully based on the CMS or e-commerce integrations. The content management system for developing Web apps is an integrated application that is used to create, deploy, manage and store content on Web pages.

Joomla is one of the key CMS platforms. It is written in PHP to enable building of websites and powerful online applications. Many aspects, including its ease-of-use and extensible nature, make Joomla the most popular Web based software development CMS. Joomla CMS is built on the model–view–controller (MVC) Web application framework, which can be used independent of the CMS.

Methods similar to Joomla's installation can be followed to install and deploy the OpenCart e-commerce platform. Using OpenCart, websites involving online financial transactions including online shopping, online bidding, payment gateways and related services, can be launched. OpenCart is pre-built in MS Azure; it can be installed in seconds and made ready for deployment. END

By: Dr Gaurav Kumar

The author is the MD of Magma Research and Consultancy Pvt Ltd, Ambala. He is associated with a number of academic institutes, where he delivers lectures and conducts technical workshops on the latest technologies and tools. You can contact him at *kumargaurav.in@gmail.com* or *www.gauravkumarindia.com.*

Rockstor

The Rockstar among NAS Solutions

This article is a tutorial on using Rockstor, particularly for setting up multi-period snapshots. The author discusses the various features of Rockstor and its installation, before moving on to cover snapshots and how to manage backups.

Rockstor is a Linux based NAS distro powered by the B-tree file system, also called BTRFS. This NAS and cloud storage platform is suitable for small to medium businesses and home users who have little IT experience, but who need several hundred gigabytes or even terabytes of data storage. The advantage of an open source DIY solution like Rockstor, compared to older NAS solutions, is that users can choose their own hardware and easily build an advanced NAS solution. Using a tailor made NAS solution offers many advantages over other options. In fact, many of these advantages form the 12 governing principles of the Rockstor project. These principles are: user empowerment, privacy, simplicity, open source, hardware agnosticism, security, ease of use, interoperability, cost effectiveness, scalability, resource-lightness and incremental development. Each of these principles (explained on the Rockstor wiki) is reflected in all aspects of the Rockstor project - from community-centric development and system requirements, to the implementation of feature sets and storage management. These principles will become clear once you deploy it for your use-case and tinker with its features.

The Rockstor NAS and cloud storage platform can be managed within a LAN or over the Web using a simple

Figure 1: Rockstor dashboard screenshot

and intuitive UI. Unlike several other NAS distros like FreeNAS, Open Media Vault, NAS for free, or even NAS boxes supporting BTRFS such as Synology or Netgear Ready NAS, Rockstor is the only distro that is purely based on Linux (CentOS) and BTRFS, and is fully open source under GPLv2. Also, compared to available public cloud storage options such as Google Drive, Dropbox, etc, storing data on Rockstor is safe, secure and cost-effective. Your data lives on your devices and on your premises, and you can easily expand or reduce storage capacity by adding or removing disks from the Web UI. You can even create a hybrid workflow, combining Rockstor with public cloud storage to have versatile access to your data.

Advanced Rockstor features

Rockstor makes it easy to manage advanced BTRFS features. It offers many advantages, such as mixing and matching drives of different types and sizes (you can combine your HDDs with USBs, external hard drives, etc), adding and removing disks, and easily switching between different RAID profiles.

Rockstor builds on top of Copy on Write (CoW) features of BTRFS to provide capabilities such as Snapshots (on demand or on schedule), Clones, appliance-to-appliance replication, to name a few. Snapshots can be conveniently created, cloned, scheduled and accessed by clients using NFS or Samba. Appliance management tasks such as appliance-to-appliance replication, adding a new appliance and making appliances aware of each other can be easily done.

All of the above operations can be carried out easily over the Web UI. So, with the combination of the Web UI and the command line, users with little IT experience as well as advanced users, and everyone else in between, can easily use Rockstor.

Installation

Installing Rockstor is straightforward and very similar to CentOS or Fedora. The Quickstart guide, available with Rockstor documentation, is a good place to start. The website (*http://rockstor.com*) gives the download instructions. For this tutorial, I have used my production Rockstor system that runs on a HP ProLiant Microserver with four disks, and runs Rockstor version 3.8-9. You could use Oracle VirtualBox or other hypervisors like KVM, VMware, Xen, etc, for installation. If you install Rockstor on your preferred hardware (bare metal), then you will need either a CD or a bootable USB (you can either create your own or buy one from the Rockstor shop).

The Snapshot feature of Rockstor

A snapshot is the point-in-time state of a share (folder). Snapshot technology in general, and snapshot implementation of Copy-on-Write (CoW) in BTRFS (Linux) in particular, offers users simple data backup, data protection and data recovery tools.

Some possible real world scenarios that could use Snapshot technology

Imagine you are working on an important project such as writing a book or a grant proposal, editing a video you shot during your vacation, recording a podcast, editing your photos, creating a presentation, writing code or anything else of significance. You could encounter a number of problems while doing so:

1. You create your work (say four paragraphs of a proposal or a video clip), and save it. You then edit and save it again (say, you delete a paragraph or remove some video footage). Then, for some reason, you want to compare your work and need to get back the paragraph you deleted or the video footage you removed. For your document, you may have maintained a manual version control and listed down every possible change, or most likely not done it at all. On the other hand, you might have set up a backup to be saved automatically. Retrieving your video editing project could be even more difficult.
2. You are working on your project and your computer battery suddenly dies, or your computer freezes and needs a reboot; you could suddenly lose work and in case you forgot to save it, you may have to start all over again.

There could be a number of other problems that cause you to lose your work, resulting in wasted time, effort and frustration. So, what are your options? Storage systems such as Rockstor offer the Snapshots technology that is designed to solve problems like those described above. These offer not just a solution, but also allow you to create a strategy to proactively manage these issues before they arise.

Figure 2: Disks

With the Rockstor Snapshots feature, you have the ability to create automatic point-in-time backups of your work, irrespective of the nature of the work (whether documents, videos, music, code, presentation, etc), and also build a strategy to manage backups.

Setting up the Rockstor storage platform for Snapshots and other tasks

To start using Snapshots and other features, you need to set up pools, shares, Samba export, etc. So follow these easy steps and you will be ready to use Rockstor's features. Once the Rockstor platform is installed, you can set up, configure and use it from a browser.

After the initial set-up, the next step is to create storage for your data. During installation on VirtualBox, you must have created virtual hard disks or added disks, if you are using bare metal. Now, we will go a step further and create pools and shares. A pool is essentially a collection of disks with a user defined redundancy (RAID) strategy. The size of a pool can be increased or decreased by adding or removing disks. Available RAID options are RAID 0, RAID 1, RAID 5, RAID 6 and RAID 10. And a share is essentially a chunk of storage space carved out of a pool. A share can also be thought of as a directory.

Disks

On the Rockstor Web UI, under the 'Storage' tab, you will see the disks. Figure 2 shows the disks that are available on my Rockstor box. Notice that disk 'sda3' has the system created 'rockstor_rockstor' pool containing the OS. It's recommended that this pool be left alone for the OS and not be used except for maybe small shares, if absolutely necessary. Resizing, compression or deletion of this pool is not permitted.

Pool

To create a pool using the Web UI, go to *Storage->Pools*. Click on 'Create Pools' and specify the 'Name', 'Raid Configuration' and 'Select Disks' that should be the part of this pool. Click 'Submit' to create the pool. Advanced users can also apply different compression and BTRFS mount options while creating a pool (or afterwards). One huge advantage that Rockstor offers here is the ability to resize a pool by easily adding or removing disks. Figure 3 shows how a pool can be created. You can see the pool 'data_backup' in Figure 2.

Share

Once a pool is set up, you can create a 'share' from your pool. To create a share, go to *Storage->Shares* and click 'Create Share'. Specify the 'Name', 'Pool' and 'Size', and click 'Submit'. Similar to a pool, advanced users can apply compression options.

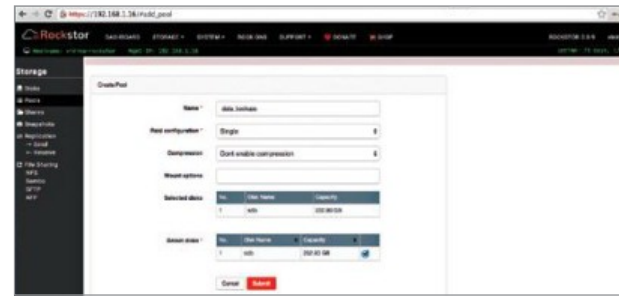

Figure 3: Pool

Figure 4: Shares

Figure 5: Create a Samba export

Internally, shares are BTRFS sub-volumes of a given filesystem (pool). A share can be exported from the Rockstor system with NFS, Samba, AFP and SFTP protocols. There is no upper limit, except available storage, on creating shares. A share can be cloned, resized and deleted from the Web UI. For this example, I have created a share called 'Marketing_Proposals' as shown in Figure 4.

To export a share using Samba, go to the Web UI, and click *File Sharing -> Samba*. Click 'Add Samba Export' and select parameters such as 'Shares to export', as well as permissions like 'Browsable', 'Read only', etc. For the 'Admin Users' field, choose the username for your Rockstor account. For example, if the account name is 'admin', select that from the drop down list. If it is your own name, choose that from the drop down list. Refer to Figure 5 for the details of my form.

Samba export of the share

Windows users can click the 'Enable Shadow Copy' option. For more details on 'Shadow Copy', follow the

Figure 6: Samba export in Mac

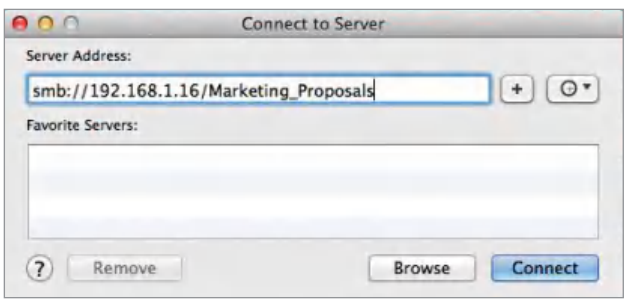

Figure 7: Creating Samba export

Figure 8: Authentication for Samba export

documentation at *http://rockstor.com/docs/windows-shadow-copy.html*. For a Mac machine (as in my case), do not click the 'Enable Shadow Copy' option. After you click 'Submit', you can mount or access the share from a client such as your laptop.

In order to mount the share, open the *Finder* and click *Go ->Connect to Server*, and type the address of the shared folder *smb://192.168.1.16/Marketing_Proposals*. Click 'Connect' and the *Finder* will pop-up the Rockstor appliance and authentication dialogue where you enter username and password (the same username and password that you are using to access the Rockstor machine). Once

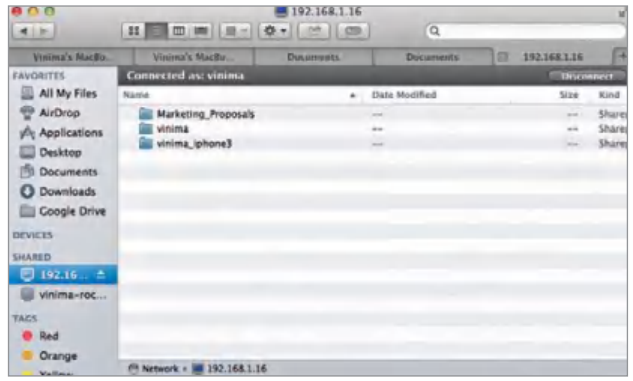

Figure 9: Mounted shares

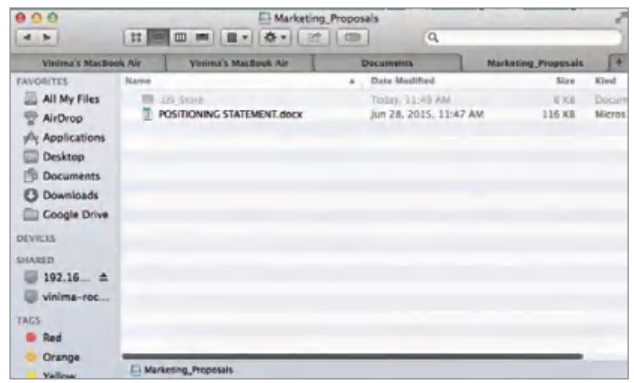

Figure 10: Content of mounted share

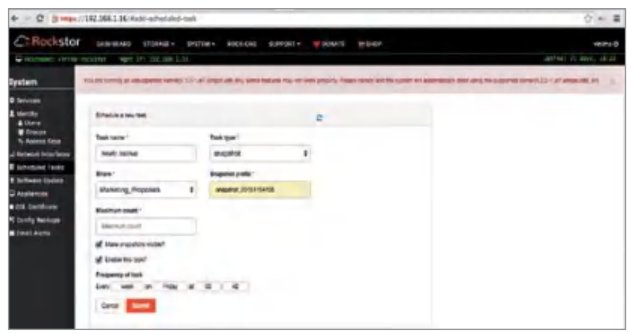

Figure 11: Scheduled snapshots

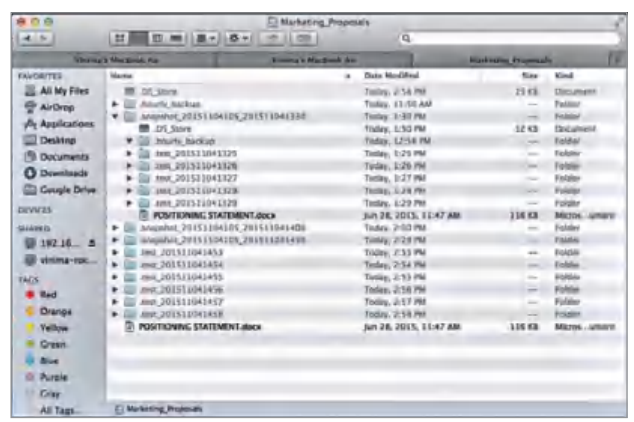

Figure 12: Snapshots on Mac

the share is mounted, copy the proposal document that you are working on from your machine to this share, and we are then ready to set up snapshots and their management strategy.

You can refer to the screenshots in Figures 6 through Figure 10 for clarity.

Setting up snapshots

For setting automatic snapshots, go to 'Scheduled Tasks' under 'System' (and click' Schedule a Task') or 'Snapshots' under 'Storage' (and click 'Schedule') of the Web UI (basically, these are two ways of getting to the same place). The images are shown in Figures 11 and 12. You can configure Rockstor to take snapshots every 30 minutes, an hour, day, week, month, year, or across multiple time periods. This way, you can recover whatever copy you had saved at a particular time.

As Snapshots are created per schedule, these will start appearing inside the share (Marketing_Proposals). Going back to our example above, if you want to retrieve the paragraph you edited and removed, and compare it to your current work, you can simply retrieve your back-up from the last 30 minutes or past hour. Figure 12 shows how snapshots are visible on your mounted folder. If you click inside a folder, you will have prior versions of your snapshot. Snapshots, therefore, ensure that you not just have the current version of your work but also a trail of all prior versions saved over a period of time. Figure 11 shows how the set-up screen for a snapshot looks.

Manage your backups using snapshots

Backups can be managed by using multi-period snapshots, a technique or strategy that allows you to schedule snapshots across multiple time periods. If you are actively working on a project for a week, you may want to back up every 30 minutes or hour of your work while you are actively working on it during that week. However, once you are done, you may want to back up every month, and then every year for archival purposes. Multi-period snapshots help you achieve this 'self-thinning' and manage your backups and storage space.

I will now describe one multi-period snapshot and self-thinning backup management strategy for the proposal document. You can modify this strategy or come up with one that fits your needs better.

While setting up a snapshot task, there is a field called *Maximum count of snapshots* where you can specify the maximum number of snapshots you want (Figure 11). Say you are setting up an hourly backup, and you want to back up every hour for seven days of a week, then you can set the maximum count to 250 (greater than 24 x 7) snapshots. Once the snapshot count reaches 250, the newer snapshots start overwriting

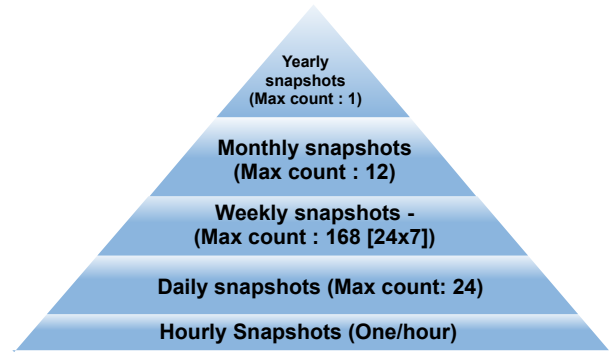

Figure 13: Self-thinning backups representation

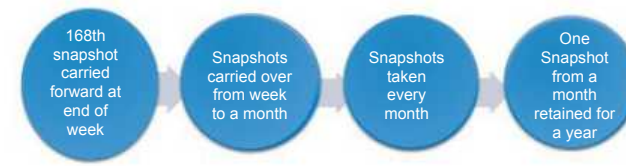

Figure 14: Multi-period snapshots carried over time

the older ones. Let us also assume that your project was a week long, and for now, you are done, but may make minor edits later. So, you want to retain the 168th snapshot, carry it forward and delete the rest. Say, you set up the snapshot tasks at 10 p.m. on a Monday, so the 168th snapshot will be taken at 9 p.m. on Sunday. You can set up a weekly snapshot task for 9 p.m. of Sunday of that week and move off the daily backups.

You could also choose to carry forward this weekly snapshot for a month. And you can choose a snapshot for a particular day (say, 9 p.m. on the third Sunday of September) and delete the rest. So, now your snapshots will be taken every month. This allows you to make minor changes to your backup and then save a copy every month.

Assuming that you have no need for a monthly backup, and now you want to retain one snapshot for a month in a year for archival purposes and delete the rest, you can do that too (See Figures 13 and 14).

This way, you can start with a lot of snapshots (starting with 168 a week) when you are actively working on a project, and then just save one copy of the final version for archiving.

Similarly, you could create your own strategy for daily, weekly, monthly and yearly snapshots and backups. END

References

[1] *http://rockstor.com.*
[2] *http://rockstor.com/docs/*
[3] *http://forum.rockstor.com/*

By: Vinima Aggarwal

The author is a technology and marketing enthusiast and describes herself as a full stack marketer. She is currently contributing to the Rockstor project and its marketing.

A First Look at Open Computer and Software Inventory–Next Generation (OCS-NG)

OCS inventory NG is cross-platform software, which supports most OSs. It is based on the popular LAMP or WAMP solutions stack. Using a standardised Web interface of HTTP protocols, the network agents communicate with the management server via XML protocols. By deploying it, the systems admin's job can be made much easier.

OCS-NG is designed to ease the systems admin's job of managing the inventory and updating information about hardware. Licensed under GPL, it is available in 11 languages. It is also available for various flavours of Linux, Windows, Solaris, HP-UNIX, AIX and Mac OS. And it even has an option called 'Remote deployment of software'. The unique feature of OCS-NG is auto rebuild. If your asset database gets corrupted and you have no backups, OCS-NG will rebuild it by re-collecting the data from agents through locally stored data.

OCS-NG consists of four major components that are listed below.

Administrative console: This can be accessed via an Internet browser, to manage the OCS-NG inventory. Usually, it will query the data from the database server.

Communication server: This manages the communication between agents and the database server via http.

Database server: Actual data is stored here. The default database is MySQL.

Deployment server: This is used to store deployment configurations and it works using https for high security.

All the components can be installed on a single server or on multiple servers. The best combination will be a DB server + a communication server and admin server + a console server. Figure 1 gives the communication methodology and architecture of OCS-NG. For high availability (HA) purposes we can use the admin server as the backup database server. If you have a small number of computers, e.g., 200+ nodes, you can deploy all components on a single server which has 4GB RAM running on Linux. If you choose multiple servers, Linux is preferred.

Let's look at the pre-requisites for each component. I tried testing OCS by installing all the modules in one CentOS 5.6 server. The following are installation/configuration examples, with a single server installation of OCSNG 2.1.2.

Preparing the servers/components

a. ***Database server:*** For MySQL 4.1 or greater with InnoDB engine active, run the following commands:

```
# yum install mysql-server
# service mysqld start
# chkconfig mysqld on
```

To check whether the InnoDB engine is active or not,

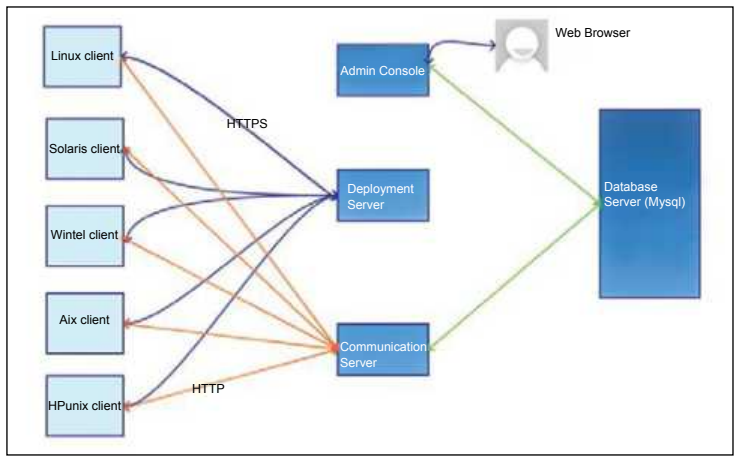

Figure 1: OCS-NG architecture

Figure 2: The complete OCS inventory

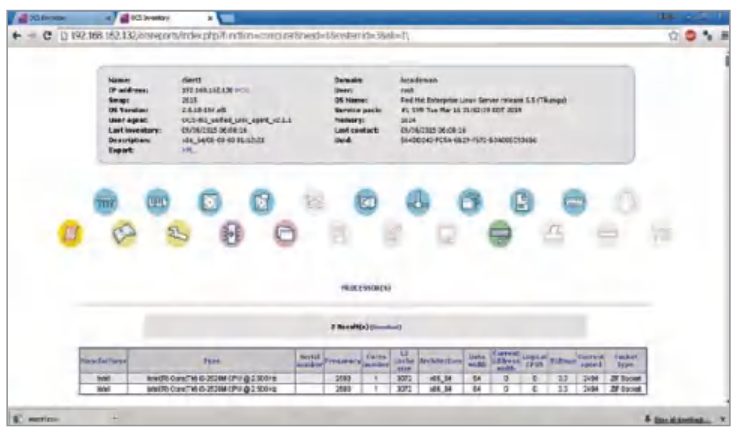

Figure 3: Inventory output

enter the following command:

```
# mysql> SHOW ENGINE INNODB STATUS
```

The default port for MySQL is 3306. Make sure your firewall or Iptables do not block Port 3306.

If you want to secure your MySQL server, run the following code:

```
#mysql_secure_installation
```

By default, the root password for MySQL will be *null*. Please set the root password immediately after MySQL is installed.

b. ***Administration console:*** What are required are Apache 2.x, Perl and PHP 4.1 or greater, along with GD and ZIP support enabled.

c. ***Deployment server:*** Apache with SSL enabled.

d. ***Communication server:*** Apache 2.x and Perl along with additional modules.

We have to install the rpms listed below. Most of the rpms are available on installation DVDs and the rest can be downloaded from the Web. I downloaded the missing rpms and created a separate repository as most of the following rpms are needed on the client side too. You can use Yum to install all the rpms. This will resolve all the dependencies automatically.

- perl-Apache2-SOAP-0.73-1.el5.noarch.rpm
- perl-Apache-DBI-1.08-1.el5.rf.noarch.rpm
- perl-Class-Inspector-1.24-1.el5.rf.noarch.rpm
- perl-DBD-MySQL-4.014-1.el5.rfx.x86_64.rpm
- perl-Digest-SHA1-2.13-1.el5.rfx.x86_64.rpm
- perl-FCGI-0.68-1.el5.rf.x86_64.rpm
- perl-IO-Socket-SSL-1.01-2.el5.noarch.rpm
- perl-Net-IP-1.25-3.el5.noarch.rpm
- perl-SOAP-Lite-0.712-3.el5.rf.noarch.rpm
- perl-Task-Weaken-1.03-1.el5.rf.noarch.rpm
- perl-version-0.91-1.el5.rf.x86_64.rpm
- perl-XML-Entities-1.0000-1.el5.rf.noarch.rpm
- perl-XML-Simple-2.14-8.0.el5.noarch.rpm
- php-mbstring-5.1.6-44.el5_10.x86_64.rpm
- php-pecl-zip-1.8.10-3.el5.x86_64.rpm
- php.x86_64 0:5.1.6-45.el5_11
- php-cli.x86_64 0:5.1.6-45.el5_11
- php-common.x86_64 0:5.1.6-45.el5_11
- php-ldap.x86_64 0:5.1.6-45.el5_11

OCS-NG is now ready to install. Download the OCS-NG server *gzip* file from *http://www.ocsinventory-ng.org/* ; copy it to the test folder and install.

Create a folder *OCSNG_UNIX_SERVER-2.1.2* and extract files under it, as follows:

```
#tar zxvf OCSNG_UNIX_SERVER-2.1.2.tar.gz
# cd OCSNG_UNIX_SERVER-2.1.2
#./setup.sh-> make sure you running this script as a root user
```

Answer the following questions during installation, as indicated below.

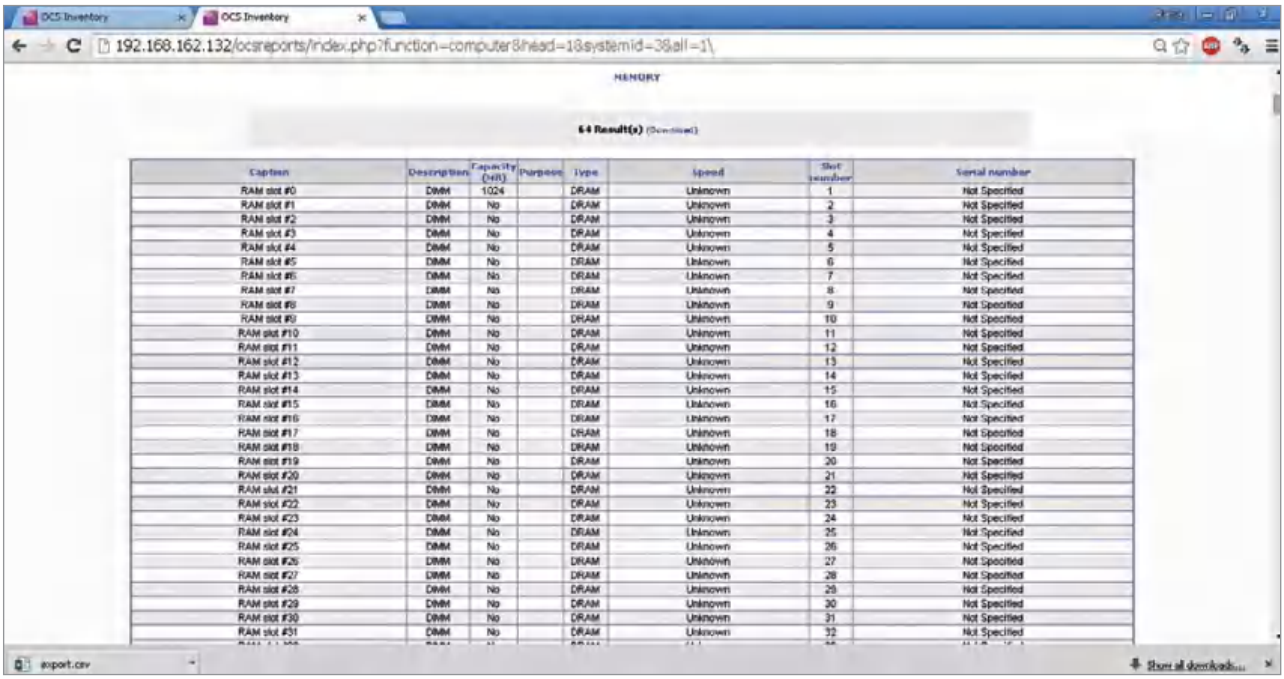

Figure 4: Inventory output

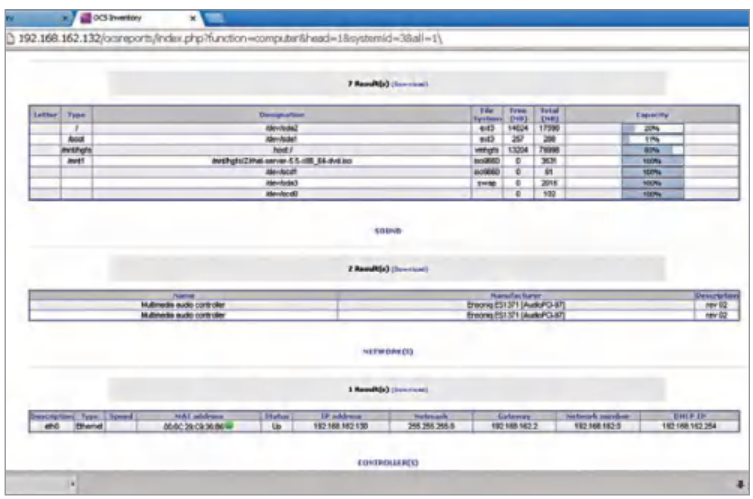

Figure 5: Inventory output

- Which host is running *database server*? It's localhost/Ipaddress.
- On which port is the *database server* running? The default port 3306.
- Where is the Apache daemon binary *[/usr/sbin/httpd]*? For default, press *enter.*
- Where is the Apache main configuration file? The default configuration file is */etc/httpd/conf/httpd.conf*, which you choose depending on your requirement.
- Which user account is running the Apache Web server? Choose Apache.
- Which user group is running the Apache Web server? Choose Apache.
- Where is Apache Include configuration directory? The default is */etc/httpd/conf.d*
- Where is the PERL Interpreter binary? The default is */usr/bin/perl*
- Do you wish to set up a communication server on this computer ([y]/n)? If you want to do so on the same one, press 'y' or 'n'. 'y' is the default here.
- Where should one put the communication server log directory? The default location is */var/log/ocsinventory-server*
- Where should one put the communication server plugins configuration files? The default location is */etc/ocsinventory-server/plugin*
- Where should the communication server plugins Perl modules files be placed? The default location is */etc/ocsinventory-server/perl*

After this, it will check all Perl modules and may throw an error for a missing SOAP::LITE module, even though you've installed it. Just continue to the next questions.

- Do you want to allow the set-up renaming the communication server's Apache configuration file to 'z-ocsinventory-server.conf' ([y]/n)? Press 'y' and then *Enter*.
- Do you wish to set up the administration server (Web Administration Console) on this computer ([y]/n)? Press 'y' and then *Enter*.
- Do you wish to continue ([y]/n)? -> You can select 'y' as this is our first installation and not an upgrade.
- Where should one copy the administration server static files for the PHP Web console? The default is */usr/share/ocsinventory-reports*
- Where should one create writable/cache directories

for deployment packages, administration console logs, IPDiscover and SNMP? The default is */var/lib/ocsinventory-reports*

The installation logs are created under */test/OCSNG_UNIX_SERVER-2.1.2/ocs_server_setup.log*. So you can store these logs for future reference.

Finally, restart Apache to reread the updated configurations.

```
#service httpd restart
```

Configuration

Connect to the admin console by typing *http://localhost/ocsreports/* in a browser. Replace localhost with your server IP if you are working remotely.

If you get an error, then use the URL *http://localhost/ocsreports/install-php* to initiate the configuration again.

Use the MySQL login name *root*, let the password field remain blank and the rest of the values be left as they are. By default, MySQL will come with a blank root password. Modify the root password immediately after configuration.

You may get the following error: 'Not able to connect to database.' The ocsweb database may not be created by default, so create it manually by using the following commands:

```
 #mysql
#show databases; -> To list DB's
#mysql> create database ocsweb;-> To create DB ocsweb
```

You can now log in to OCS to configure your inventory systems. Type *http://localhost/ocsreports/* in a browser and you should get a login screen. The default username/password is admin/admin. Now we will jump to OCS agent installation on client servers.

Agent installation/configuration

To install the OCS agent, we have to ensure that the following rpms are installed as a pre-requisite.

```
Compress::Zlib version 1.33
Digest::MD5 version 2.33
LWP::UserAgent version 5.800
Net::IP version 1.21
Net::SSLeay version 1.25
XML::Simple version 2.12
```

Download the *Ocsinventory-Unix-Agent-2.1.1.tar.gz* package, extract it and compile it by using the following code:

```
# tar zxvf Ocsinventory-Unix-Agent-2.1.1.tar.gz
# cd Ocsinventory-Unix-Agent-2.1.1
# perl Makefile.pl
# make
# make install
```

Once it is successfully compiled, it will start executing the *postinstallation.pl* script, which will ask the questions listed below. Be ready with the IP address of the master server running the OCS-NG service.

- Do you want to configure the agent? If so, enter 'y'
- Where do you want to write the configuration file? Choose the following options:
  0 -> */etc/ocsinventory*
  1 -> */usr/local/etc/ocsinventory*
  2 -> */etc/ocsinventory-agent*
- What is the address of your OCS server? IP ADDRESS of your OCS master.
- Do you need the credentials for the server? Enter 'n'.
- Do you want to install the cron task in */etc/cron.d?* If you want to run this as a cron job, press 'y'; otherwise 'n'.
- Where do you want the agent to store its files? The default location is */var/lib/ocsinventory-agent*
- Should I remove the old Linux_agent? Enter 'y'.
- Do you want to disable the SSL CA verification configuration option? Choose 'n' if you need a secured one.
- Do you want to use the OCS-Inventory software deployment feature? Enter 'y'.
- Do you want to use the OCS-Inventory SNMP scans feature? Enter 'y'.
- Do you want to send an inventory of this machine? Enter 'y'.

Finally, you will get the following output, which will show successful completion:

```
Creating /var/lib/ocsinventory-agent/http:__192.168.162.133_
ocsinventory/snmp directory...
Copying SNMP MIBs XML files...
Activating modules if needed...
Launching OCS Inventory NG Unix Unified agent...
 -> Success!
```

We have completed the agent side configuration too. Let's jump to the OCS admin console *http://localhost/ocsreports/*. Run the inventory agent from the client side, and it will automatically register in the OCS database. If it is not visible in the OCS admin console, please check the agent installation logs.

Let us view the icons shown in the OCS intial screen given in Figure 6. A brief description of each icon is given below:

1. Views all computers - once you click on it, you will get a screen like the one shown in Figure 3.
2. Views computers group-wise. You can create your own group like Prod/QA/Dev.
3. Views the computers based on your tag. For example, you can add or customise a tag (keyword) to all computers. If I add the city of the data centre as a tag, then I can search the servers in Chennai city.
4. Lists all the software installed on the server.

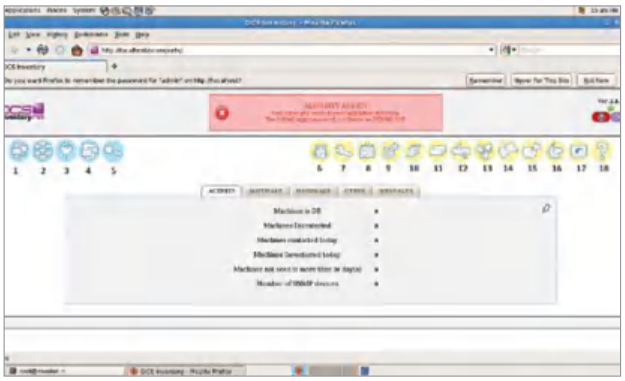

Figure 6: OCS initial screen

5. This is to search computers using a name, like, IP Addr, MAC addr, cpu, architecture and many more. And you can use multiple operators to compare groups.
6. This deploys and activates packages/apps. First, you have to build the package and deploy it on client servers. The important attributes for building a package bundle are:
   a) Priority: Level 0 to 9. 9 is the least and 0 is the highest priority. This will be used to decide which package has to be installed first when deploying multiple packages.
   b) File: Zip for window and tar.gz for Linux.
   c) Action – Store: Just download the bundle at the client server.
   d) Execute: This will deploy and execute the installer file from the bundle. Either *.exe* for Windows or an rpm in the case of Linux. It will just execute the command from the command box and won't return anything like 'success' or 'failure'.
   e) Launch: This is to deploy a bundle with the command and it will return the value. Here you can use your own script to deploy patches.
   f) You can configure the user notification if you want to notify users.

   Once done, you can activate the bundle to make it ready for deployment.
7. Config: You can configure the default attributes related to the server, IPDiscovery. We will discuss a few important attributes below:
   a) Frequency: This refers to the frequency (in days) to collect inventory.
   b) Inventory_Diff: This is for whether to enable differential inventory or not (like differential backup).
   c) Inventory Transaction: If this is enabled, it will not store inventory information in case of any partial error.
   d) IPDiscover: This is to find the number of agents running IPDiscovery per subnet.
   e) Deploy: This is for whether to enable automatic deployment or not.
   f) Download_timeout: This is the maximum time a package can take to download to the client. If it exceeds the limit, cancel the download and log error message.
   g) Local_port: This is the TCP port of the OCS communication server.
   h) Local_server: This is the communication server's name.
   i) Loglevel: This is the level of log creation. */var/log/ocsinventory-NG* for Linux and *<appfolder>\xampp\apache\logs* for Windows.
   j) LDAP configuration: This is to configure LDAP credentials.
8. IP Discover: This is to auto discover the devices in a network or subnet. This feature will auto discover servers in a network.
9. Registry: This modifies Windows' registry values.
10. Administrative data: This stores custom data for each server. For example, you may store the data centre's name, rack number, etc, by creating custom values here.
11. Duplicates: This checks duplicates of IP, hostname and mac address, or a combination of all these parameters.
12. Dictionary: This categorises the detected packages manually.
    a. New: Uncategorised softwares/packages
    b. Ignored: Cannot import GLPI under these packages
    c. Unchanged: The packages under this category will be imported, as is
13. Plugins: This lists installed plugins and the management of these plugins.
14. Logs: This views the logs and imports them as *.csv* files too.
15. Statistics: This gets the statistics about valid connections and failed connections, as well as successful and errored deployments.
16. Users: This creates users with different roles like Super admin/Admin/requester.
    a. Super admin: Can change the entire configuration
    b. Admin: Can add/modify users and has fewer controls
    c. Requester: Has no modification access
17. Local import: Another good feature in OCS is that you can maintain inventory for non-networked servers also. You can run the inventory agent and export the file as a *.ocs* file from source, and import the same in the admin console using this option. So you can include servers in DMZ without actually connecting to it.
18. Help: Shows the wiki page for more help.

The list of options is very long. We can see what information we get from the inventory through the screenshots given in Figures 4, 5, and 6.

The inventory will list all the details of the installed hardware and software as shown in these figures. There are many customisations which can be done in OCS. Remember that these are only Linux agents. END

By: Panchavarnam Ramasamy

The author is a software architect at Mphasis. He has been in the IT industry since 1997, and is proficient in UNIX and cloud technologies. He is passionate about researching and testing open source tools. He can be reached at *sparcrams@yahoo.co.in*

The Basics of Graphic Network Simulator 3

GNS3 is a software emulator for networks. The GNS3 website describes it as the '…software that empowers network professionals.' This article provides the reader with the basics of GNS3 and a few test labs, which could be tried out by interested readers.

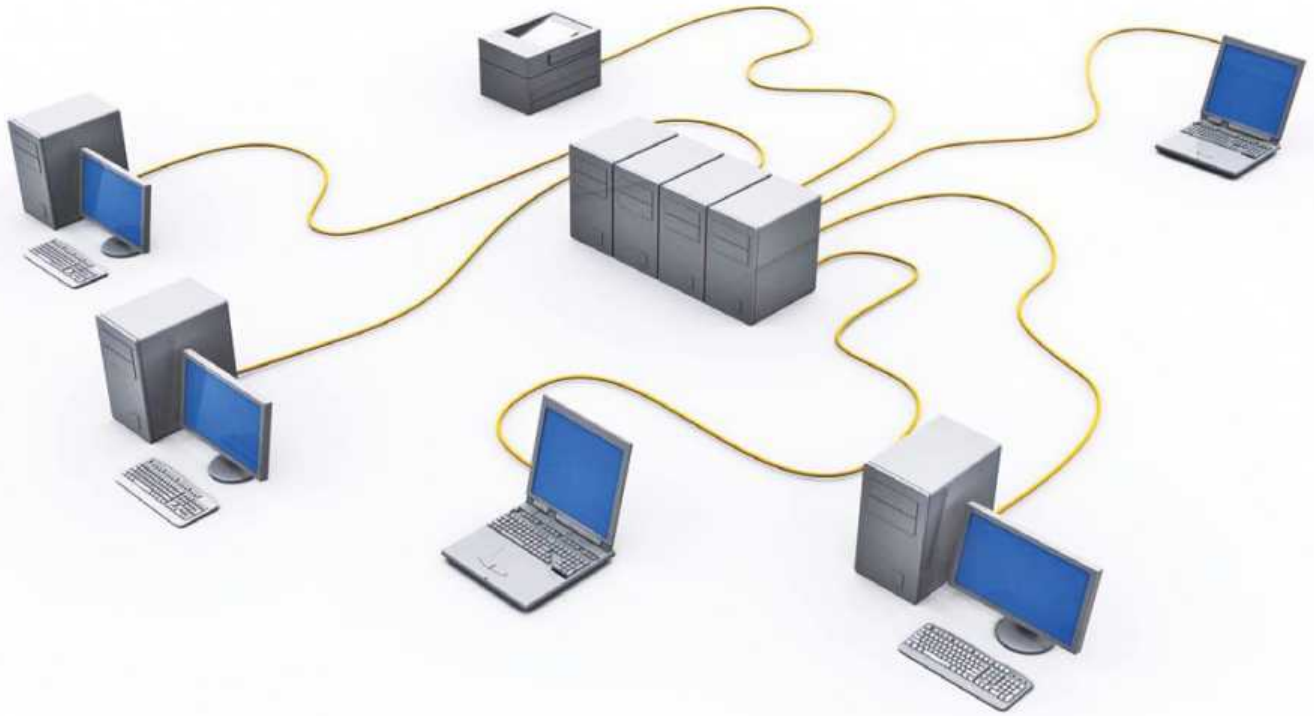

Reconfiguring a live network - for example, while implementing access control lists (ACLs) in an enterprise network - can be a challenging task. Troubleshooting such networks could also be really difficult if strict change management policies are in place, and specially if downtime and disruption are to be avoided under any circumstances.

Is there a way out of this? Yes! Many professionals use GNS3 to simulate complex networks. It is also used by students to prepare for lab tests conducted by various certification providers.

History

In 2005, Christophe Fillot wrote Dynamips - a Cisco router simulator. It had a text based interface and could simulate only one router on one PC, at a time. Version 0.2.5 of Dynamips, released in 2006, introduced hypervisor mode and enabled multiple router simulations on one PC. It also provided a simple virtual switch.

GNS3 was introduced in September 2007 with a GUI, drag-and-drop facilities for adding network components, and easy storage of created objects and topology.

gns3.com describes GNS3 as 'The software that empowers network professionals.'

Dependencies

GNS3 depends on several libraries and components. Successfully installing these dependencies is a prerequisite. The important dependencies are listed below.

- Dynamips: Base of GNS3, a Cisco router simulator.
- Dynagen: A text based front-end to Dynamips, which uses the hypervisor mode to communicate with Dynamips.

Figure 1: GNS3 starting

Figure 2: Adding new IOS router

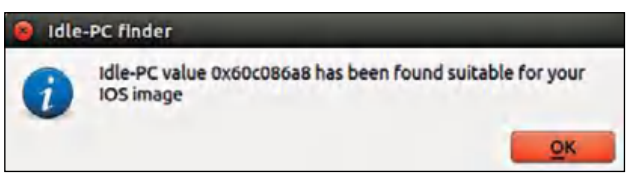

Figure 3: Suitable Idle-PC value

It uses a simple configuration file for specifying virtual router configurations, and enables interconnecting of routers and WAN technologies such as ATM, frame relays and switches. Further, it provides management CLI for device functions such as start, stop, suspend, reload, console connection, etc.

- Python: Dynagen is written in Python.
- VPCS: This is a virtual PC simulator, which can simulate up to nine PCs. It is possible to use the *ping* and *traceroute* commands from or towards these virtual PCs. Remember that VPCS is just a program running on Linux or Windows supporting a few network commands. But, it becomes difficult to configure a network without VPCS. Earlier, VPCS was not available, and users were required to configure an additional router in place of a host (PC) for testing.
- QEMU: A generic and open source machine emulator and virtualiser. As a machine emulator, it enables programs made for one system to be run on different machines (e.g., programs developed for ARM can be run on a PC).

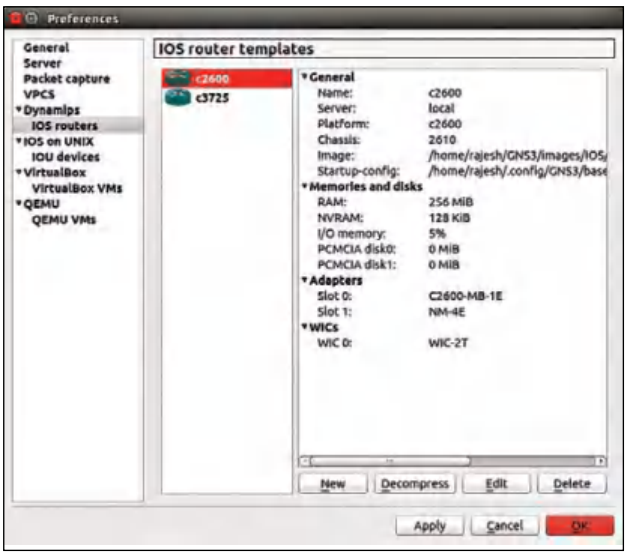

Figure 4: IOS router templates

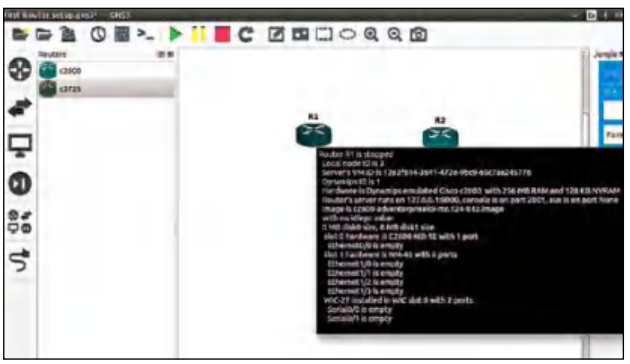

Figure 5: Drag-and-drop router

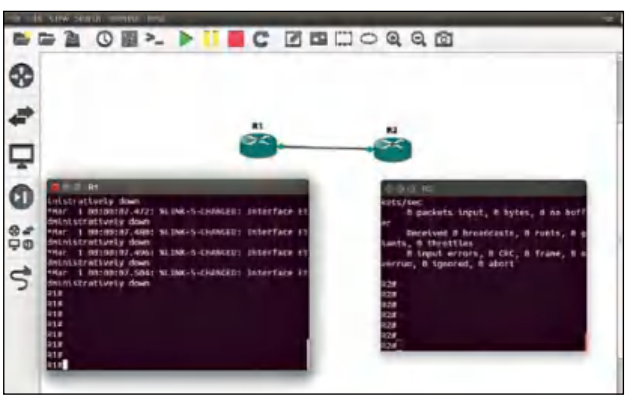

Figure 6: Console for two routers established

As a virtualiser, it achieves near native performances by executing the guest code directly on the host CPU.

- cpulimit: This attempts to limit the real CPU usage of a process. It is helpful to limit CPU usage of devices created in GNS3.
- VirtualBox: A virtualiser for x86 hardware, targeted at server, desktop and embedded use—a cross-platform virtualisation application.

- IOUYAP: This provides network support for IOU (IOS for UNIX).

Supported images

GNS3 (Dynamips) uses actual device IOS images to run the actual IOS rather than simulating device commands. An example of a device simulator is Cisco Packet Tracer. Only selected versions of IOS, and selected modules, cards, processing engines, port adapters, etc, are supported. A list of images supported by GNS3 is given below. Please remember that only selected devices from the mentioned device series are supported.

Cisco router series	1700, 2600, 3600, 3700, 7200	**Cisco PIX firewall**	PIX 525
Cisco ASA firewall	ASA 5520	**Cisco IDS sensor**	IDS 4235/4215
Cisco IOU	IOU is IOS for UNIX. These images can be used with GNS3 for switching.	**Juniper router**	JunOS series M (JunOS of the Juniper router is based on FreeBSD)

Detailed list of supported hardware

A detailed list of supported hardware is available at *https://community.gns3.com/docs/DOC-1708*.

An interesting guideline mentioned in this link is to use older images supporting your configuration requirement, for better speed.

System requirements

To run real IOS image(s) in virtual environment, CPU and RAM resources from your computer are required. Virtualisation also requires additional RAM and CPU capacity.

Considering these requirements, your computer system should have CPU and RAM resources sufficient for the desired IOS image and network topology.

If you wish to configure complex topologies, the best choice will be a 64-bit Linux OS, with a minimum of 4GB RAM, or 64-bit Windows with a minimum of 8GB RAM. This is because, as you must be aware, all 32-bit operating systems have a total addressable space barrier, which limits the maximum RAM limit to 3GB.

Downloading GNS3

Register with *gns3.com* and proceed to the download section. You may also use the GitHub link *https://github.com/GNS3/gns3-gui/releases*, where you can select the desired version and go ahead. The approximate download size for GNS3 1.3.7 is 65MB for Windows and 40MB for Linux.

Installation instructions

Windows: Download installables and double click.

Ubuntu 14.04: Using the command line, you need to install Python, Dynamips, GNS3 Server, GNS3 GUI, IOU dependencies, IOUYAP, VPCS, Virtual Box, QEMU and cpulimit. You may also consider installing Wireshark to capture packets for diagnostics and troubleshooting.

Step-by-step instructions for installation of GNS3 on Ubuntu 14.04 are available at *https://community.gns3.com/thread/5471*. Complete the installation. GNS3 runs in non-administrative mode, so use the command *gns3*.

Practical uses

GNS3 is a great help for students aspiring for various Cisco certifications. Preparation for most of the laboratory tests for CCNA, CCNP and even CCIE is possible using GNS3.

GNS3 is also a handy tool for professionals. Consider an enterprise with a data centre (DC), disaster recovery (DR) site, primary and redundant links, routers connecting branches to the DC and DR, etc. An incorrect change made to any of the networking devices may lead to incorrect configuration or network outage. Here, GNS3 can be used to configure the virtual network mirroring the actual set-up. This virtual set-up can be used as a test environment to make the desired changes, and then these tested changes can be applied to the live (production) network.

Did you guess an interesting dependency of GNS3, not mentioned above?

Where can Cisco IOS and IOU images be obtained? Are they provided on the GNS3 website? Can anybody provide them?

Let me make it very clear – GNS3 users will have to arrange IOS images on their own. IOS and IOU images are the property of Cisco, and so user licences are available from the company. You have been warned!

Your first lab: GNS3 1.3.7 installed on Ubuntu 14.04

Start GNS3 from the command line as follows:

```
$ gns3
```

During startup, GNS3 will check and inform you if any new updates are available.

You must follow these steps to add your first two routers:

1. Enter the name of the new project. The default project folder is */home/username/GNS3/projects/*, which can be changed from *Edit - Preferences*. Give your project a meaningful name.
2. Add router images to GNS3.
   Go to *Edit - Preferences - Dynampis - IOS Routers -*

Browse - select the desired image file, and then expand the file as prompted.

For this lab, Cisco 3725 and 2600 IOS images are used.

3. Populate the 3725 router with supported cards. The default card is GT96100-FE, a two-port 10/100 Fast Ethernet. In addition, the following models are supported.

a. Network adapters (any two of the following):
   - NM-1FE-TX: One fast Ethernet (10/100)
   - NM-4t: 4-port asynchronous (serial) module
   - NM16-ESW: 16-port 10/100 EtherSwitch modules for Cisco ISR series

b. WIC modules (any three of the following):
   - WIC-1T
   - WIC-2T

4. Set the idle PC value. When a router is just started, CPU utilisation of the computer system may rise to 100 per cent and the system may become very slow. To avoid this, the *idlepc* setting is necessary. 'Select Idle-PC value' displays the default value. Use the 'Idle-PC finder' button to check whether it is suitable. Often, the default idle PC value is suitable.
5. Save the configurations by clicking *OK* and proceed further.
6. Follow Steps 2 through to 5 to add a 2600 series router image.
7. Configure the topology. Drag and drop as many routers you wish to use in the topology, and use the *Add a Link* button to connect them by cable. This test lab has the 2600 Router R1 Ethernet0/0 to 3725 router R2 FastEthernet0/0.

Note that when you move the cursor on the router, its status and configuration is displayed in the adjacent black window.

8. Use topology devices.

- Click on the *Start / Resume all devices* (the green > button on top) to start all devices.
- Click on the >_ button to trigger console connection to all devices.
- Click on the red square button to stop routers.

Do not forget to save your project by using the *File – save project* menu. Go ahead and configure these routers for your test labs.

Changing the configuration of these router images including network/WIC modules, RAM and Flash is possible from the *Edit* menu. You can even change the base MAC ID of the router under the *Advanced* button.

Keep a watch on this column for updates in exciting world of GNS3. END

By: Rajesh Deodhar

The author is an IS auditor, network security consultant and trainer with 25+ years industry experience. He is an industrial electronics engineer with CISA, CISSP and DCL certification. Please feel free to contact him on *rajesh at omegasystems dot co dot in*.

Run Multiple Operating Systems on a Single Machine with VirtualBox

Here's a complete overview of the open source tool VirtualBox that will help you to run multiple operating systems on a single machine. This article is targeted at all those who run multiple operating systems, or would like to try out new operating systems and configure them without disturbing the existing set-up of their machines. For noobs, the article contains a hands-on demo for setting up a virtual machine (VM).

Virtualisation in computer science terms means creating virtual and isolated instances of computer resources like storage, memory, processors and memory. Virtualisation exists at various levels such as at the hardware level, operating system level and application level.

Introducing VirtualBox

In very basic terms, VirtualBox is a cross-platform virtualisation manager application, which enables the user to experience various operating systems simultaneously on Intel and AMD based platforms.

Where to get it

VirtualBox is free and open source software available for free download. One just needs to start typing the word VirtualBox on a search engine and there will be thousands of links to download from, for free. Two of these are: *http://www.filehippo.com/download_virtualbox/* and *https://www.virtualbox.org/wiki/Downloads*

Installing VirtualBox on Windows

You can install VirtualBox the same way you install any other software. Once the installer is downloaded, just double-click *VirtualBox-version-win.exe* and start the installation. The installer will start and it will provide some details with regard to installation location, type of install, and the features it will install in the next few screens. A new user can safely keep going… *Next->Next* and *Next*.

Towards the end of the installation, the installer will ask for permissions to install various device drivers; choose 'Install' and finish the installation.

Additional drivers required during installation

When you install VirtualBox, you will require a few additional drivers for your virtual machines (VMs) to work and access all the hardware attached to the host.

In Windows, as soon as you start the installation on the second screen (Figure 1), you will see a list of features that the installer will install along with the set-up. Towards the end of the installation, these are the additional packages for which set-up will ask your permission.

Let's take a brief look at what these special drivers are for.

VboxUSB: This module will enable USB support inside virtual machines, and the user will be able to attach the host's USB devices to virtual machines.

VboxNetwork: This package will install two required network drivers — *vboxnetfit* and *vboxnetapp*. Both these modules are responsible for enabling networking support to the VM.

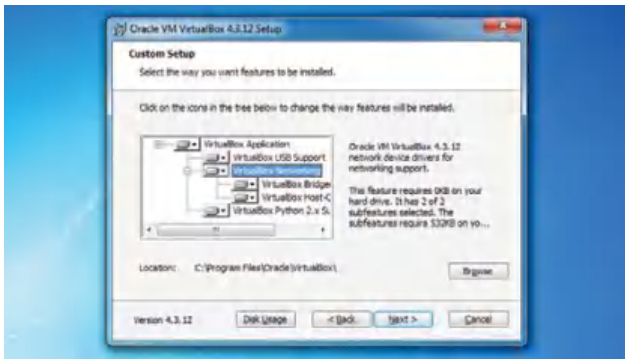

Figure 1: Windows-VB2

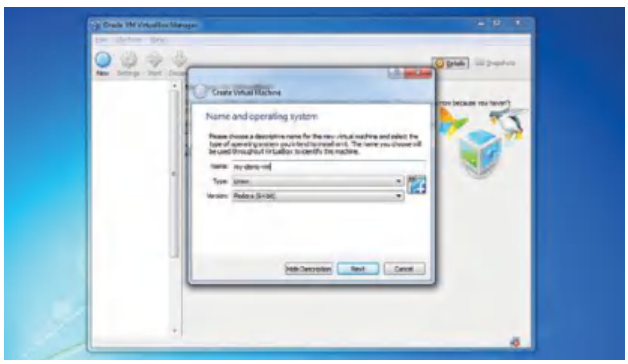

Figure 2: VM-creation2

Figure 3: VM-creation3

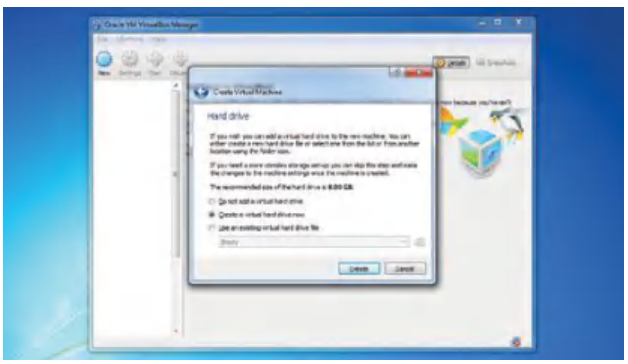

Figure 4: VM-creation4

VBox Python: This module will enable Python support to VMs.

Steps for creating your first VM using VirtualBox

Now that we have VirtualBox installed on our system, let's try and create a VM using it. Since VirtualBox is a platform-independent application, the user interface remains the same across all platforms – Windows, Linux or iOS.

The steps to create and configure a VM using VirtualBox follow. Start the VirtualBox application from the *Start* menu. You will see a welcome screen with the menu option *New*.

Click on *New* to start a new VM creation process.

1. In the next screen (Figure 2), the user will be asked to give a name, select the VM type (Windows, Linux, iOS, etc) and inform what flavour of the selected type is to be installed. For example, I have chosen Linux as my VM type and Fedora-64 bit as my flavour. I also named my VM as *my-demo-vm*.
2. Once done with the above selection, click *Next*. The next screen (Figure 3) will ask you to select how much memory you want to dedicate to your VM. The dialogue box will show a range from a minimum permitted memory range to the maximum physical memory you have on the system. I have chosen 1024MB (1GB) for my demo VM.
3. Click *Next* and the next dialogue box (Figure 4) will ask how much of disk space you want to allocate to the VM. The dialogue box will give you a few options with regard to disk selection:
   a. **Do not add a virtual hard drive:** This is in case you do not want to have any local storage to your VM. Then where should we install this? Let's leave this option for now.
   b. **Create a virtual hard drive now:** This is the default selected option and most common too.
   c. **Use an existing virtual hard drive:** This option is useful when you've already created a VM disk earlier and you want to re-use it.

   For our example, let's go with option 'b', the default, and click on *Create*.
4. The next dialogue box (Figure 5) will show a list of the types of disk drive images that the user can create. Each adds a specific feature but for a normal user, we can go with the default option again and click *Next.*
5. The next dialogue box (Figure 6) is very important with regard to hard disk drive configuration. Here, the user will have two options to choose from for hard disk space. Let's look at each of them:
   d. **Dynamically allocated:** This option will create a disk of the minimum required size and then will keep increasing it as needed, up to the maximum disk size limit allocated; but once increased, the reverse doesn't happen.
   e. **Fixed size:** This option will create the virtual disk drive of the maximum size and the same disk space will be isolated from the host's use.

INVEST. IN. INDIA.

INDIA ELECTRONICS WEEK

January 11th-13th, 2016 • Bengaluru

Six Events. Coming Together

Electronics For You expo
Innovate. Design. Manufacture. Source.

An expo for OEMs to meet electronics related suppliers and service providers

electronics Rocks

India's largest confluence of technology decision makers across the electronics industry

TEST & Measurement India

Asia's leading exposition of test & measurement products and services

LEDASIA.in

India's first event to focus on the latest technologies for LED lighting

iotshow.in

India's first large scale IoT event with an exposition and conferences

raksha India

India's first event related to the use of electronics in defence, aerospace & security

What's Happening at **IEW 2016**

- Buyer Seller Meet
- ELCINA CEO Summit
- Vendor Development Talks
- International Exhibitors
- T&M showcase
- Innovators' Zone
- ELCOMA's LED Lighting Summit
- LED technical conference
- VCs meet IoT Startups
- Technical conference for Developers and Systems Integrators
- Technical conference for Design Engineers
- Success Stories of IoT implementations

Technical Conference for electronics design engineers and embedded systems' developers

January 11th-13th, 2016 • Bengaluru

For more information, talk to us at +91-11-40596605 or email at growmybiz@efy.in

Gold Partner

KEYSIGHT TECHNOLOGIES | ARTS

Associate Partners

MOUSER ELECTRONICS | messung | erfi | Aplab | JUKI | emst | DRIVE

Tea/Coffee Partner

Supporting Association Partners

ELCINA | CEAMA | ESC | SCGI | COWE | IESA | IPC INDIA | SMTA | TEMA | IEEE

International Partners

CEC | CEIEC | TAITRA

Community Partners

ICC SALES | DIY Hacking | IoTSLR

Media Partners

Online

electronicsforu | EFY Times | OpenSource | eziefinder | EXPLARA | 10Times | LED Bazaar

Print

ELECTRONICS BAZAAR | LED | OpenSource | 3D ELECTRONICS WORLD

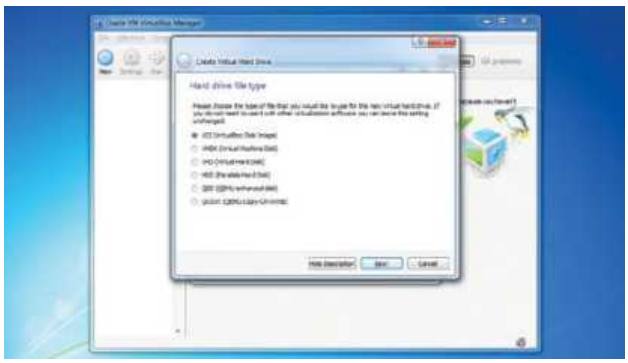

Figure 5: VM-creation5

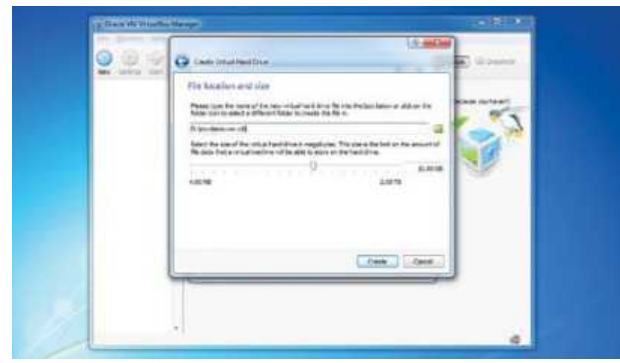

Figure 7: VM-creation7

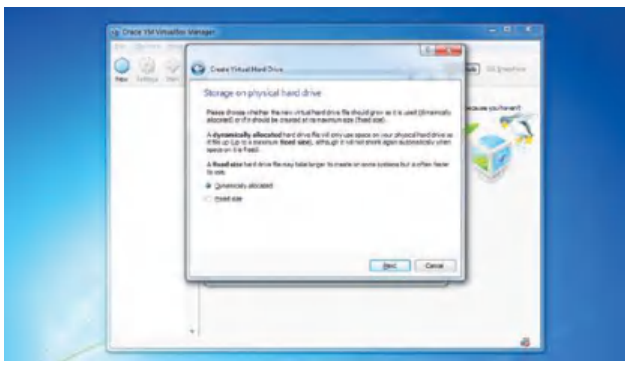

Figure 6: VM-creation6

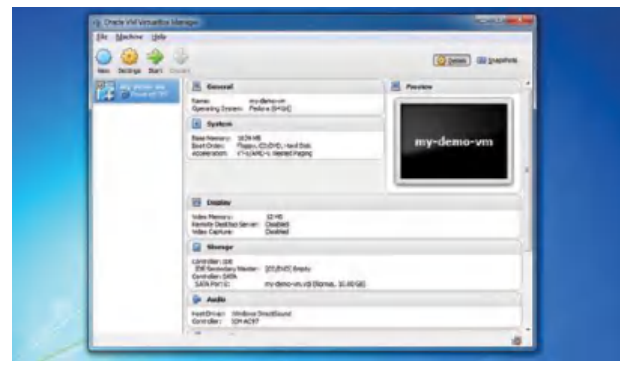

Figure 8: VM-creation8

A dynamically allocated disk allows maximum use of the host's disk space and prevents wastage of unused disk space within the VM but, on the other hand, since disk space is being allocated dynamically, the I/O operations on the VM take a hit. The decision on which of the two options to choose lies entirely with the user and the purpose the VM is being created for.

For our example, I will go with the default option, *Dynamically allocated.*

6. The next dialogue box (Figure 7) will allow you to select the virtual hard disk size, its name and also the file system location where the disk image is to be created. All choices available in the dialogue box will have their default values, but I would personally advise users to alter at least the location for disk image creation. Let's look at all the options.
   f. **Virtual hard disk driver location:** By default, this location is */user/Home/Virtual/* but I would personally suggest that you keep to a different drive so that you can maintain a better and separate track of the virtual disk drives you have created. Also, if anything nasty happens to the VM, it will not have any effect on your *C* or *User* directory.
   g. **Size of the virtual disk drive:** Users can select the desired size for the virtual disk drive.

   Once everything is configured as per desired values, you can click *Create* to actually create the disk drive image and the VM itself.
7. Our VM is now created and the VirtualBox home page will list it (Figure 8) along with the details of its basic properties.
8. It's now time to start using our VM. But wait! We haven't installed any operating system on it.
9. In the VM context, installing an OS is considered to be an operation on the VM rather than a part of the VM creation process.

Configuring a VM

In this section, we will look at how we can adjust the various properties of our VM, add or remove hardware resources, and install an operating system. We will also look at how to start using a VM as an independent OS.

1. **Configuring VM properties:** If you think that you would like to change some properties or resources allocated to the VM that were missed out during VM creation, go to the VirtualBox main page and select the 'Settings' option (Figure 9) to access all VM properties and change them, as desired. A point to note here is that you will not be able to change each and every value, but those restrictions are self-explanatory.
2. **Installing an OS on the VM:** It's now time to install the operating system on the VM. Let's look at the process in the following steps.
   a. First and foremost, we need an installation source (DVD/network boot location). For this demo I am using the DVD method and leaving the network boot for readers to try out.
   b. To associate a DVD with the VM, go to the VirtualBox main page and click on *Settings*.

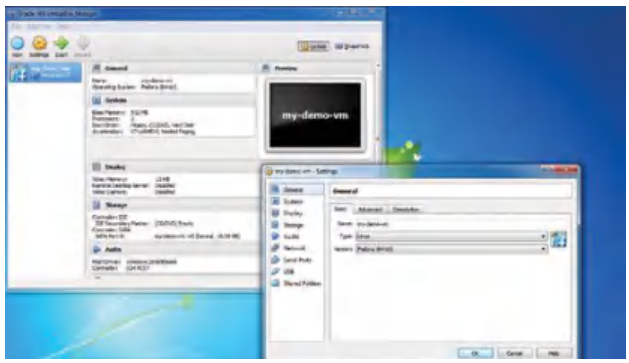

Figure 9: VM-creation9

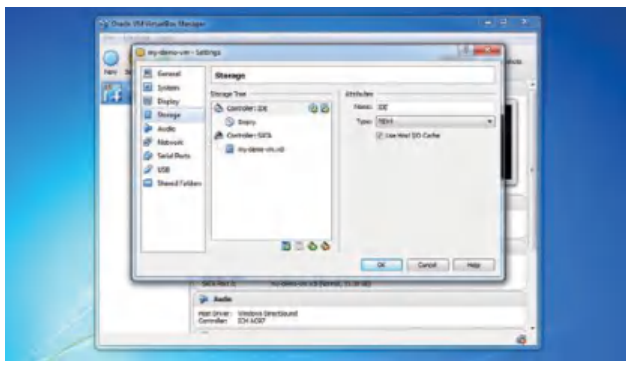

Figure 10: VM-creation10

c. On the next window, from among all the options on the left hand side, click on *Storage*.

d. On the right hand side (see Figure 10), you will see a storage tree detailing storage controllers and devices attached to the VM.

e. On the same window, the first in the tree is the Controller IDE, for which the device list is empty; this means there is no DVD attached to the VM yet.

f. Select *Controller IDE* and click on the first 'plus' icon, which should say *Add CD/DVD Device* when you scroll the mouse over it.

g. The next window will show a box, which asks if you want to add a DVD disk. Click on *Choose disk*, browse to the DVD image location and select the image as desired. The selected DVD image should show up on the storage tab under Controller IDE (Figure 11).

h. Finally, go to the *System* tab in the *Settings* menu and arrange the boot order accordingly.

i. Now, in the *Settings* window, click *OK* to save all the changes. Once all the changes are saved, you will be brought back to the VirtualBox main page.

3. And yes, we are done… Now just click the *Start* button in the VirtualBox main page and start the VM.
4. From here on, the VM will boot into the selected boot device, and you can continue installing the OS and then boot into it.

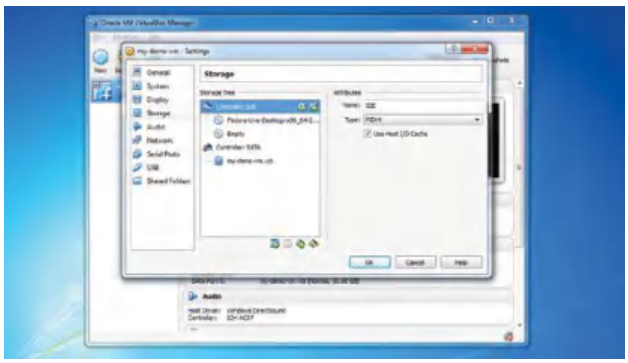

Figure 11: VM-creation12

Advanced features of VirtualBox

Now that we have looked at the basics of VirtualBox, let's look at a few advanced features. Most of these features are useful for professional users and those who manage multiple VMs across many hosts. Since it's not possible to introduce all the features that VirtualBox offers, I will try to list a few.

1. **Manageability:** VBoxManage is a command line interface to VirtualBox, and it allows you to do much more than what can be done from a GUI. More details of how to use it can be found at *https://www.virtualbox.org/manual/ch08.html*.
2. **VM groups:** The VirtualBox manager enables you to configure a group of VMs in order to manage and configure them collectively. More details are available at *https://www.virtualbox.org/manual/ch01.html#gui-vmgroups*
3. **Create snapshots:** If you are familiar with virtualisation technology, you will know what a snapshot means. VirtualBox enables you to create snapshots of the VMs and restore them back to the previous state. More details are available at *https://www.virtualbox.org/manual/ch01.html#snapshots*
4. **Creating a clone:** VirtualBox enables you to create the exact copy of an installed VM. This property is called cloning of a VM. More details are available at *https://www.virtualbox.org/manual/ch01.html#clone*

VirtualBox is one of the best VM managers available even when compared with proprietary solutions. It offers a lot more than what desktop users may require and allows users to experience a smooth multi-OS environment. END

References

[1] *https://www.virtualbox.org*
[2] *https://www.virtualbox.org/manual/UserManual.html*
[3] *http://www.pcmag.com/encyclopedia/term/53961/virtualization*
[4] *http://www.webopedia.com/TERM/V/virtualization.html*
[5] *http://www.google.com*

By: Ashish Bunkar

The author has over five years' experience as a Linux professional. He likes sharing his interests in SAN and virtualisation technologies.

Configuring Chef Server to Monitor Chef Clients

The previous article in DevOps Corner gave an overview of Chef and the procedure to set up a basic sandbox environment with Chef Server. In this article, the authors look at how to install and configure client VMs and Chef agents on them. Subsequently, they move on to explore how Chef agents connect to Chef Server, and how to configure agents to enable monitoring the clients from the server.

So far, we have learnt how to set up Chef Server and configure it by running the *Chef-server-ctl* reconfigure command. We have also learnt how to verify the server configuration by running the test command. Now, the server is almost ready to be used, though there's still some work to be done to get it started and communicating with agents. For this, we need to create an administrator user and organisation. The purpose of creating an administrator user is to control the Chef environment (server and agent), and of creating an organisation is to set up a profile for the Chef environment.

The syntax for a command to create an admin user is as follows:

```
#Chef-server-ctl user-create<username><first name><last name><email><password>--filename<full path of pem file>   --- here pem is the private key file created in the specified folder.
For example:
# Chef-server-ctl user-create magesh magesh kasthuri magesh.kasthuri@wipro.com password --filename /root/.Chef/magesh.pem
```

Then, we can run *org-create* to create the profile for the Chef environment or organisation. The syntax for the command is as follows:

```
# Chef-server-ctl org-create<Org name><Org description>–association-user<Admin user associated with this org>
For example:
# Chef-server-ctl org-create magsandbox "Mag-Sanbox Setup" –association-user magesh
```

After creating the organisation, we can run the test command again to verify if the user is associated properly and the server can start with the organisation profile.

The Chef management console is a Web based tool which provides a graphical user interface (GUI) to manage environments, user roles, cookbooks, etc. After verifying the Chef Server set-up, we need to install the management console by executing the following commands:

```
# Chef-server-ctl install opscode-manage
# Chef-server-ctl reconfigure
```

The command to configure and bring up the management console is:

```
# opscode-manage-ctl reconfigure
```

There are two more steps to complete the server installation and one of them is to install the push job service. This is called the *opscode-push-jobs-server* and it helps to push jobs or actions to the Chef agent machines. Once this is pushed, the agent will pick them up and run the job or action in the respective machines. This is shown in Figure 2. In this screenshot, you will notice a service called *opscode-reporting*, which serves reporting data to the management console to enable the user to visualise the operations from the server and agents.

Installing *opscode-reporting* is the last step in setting up Chef Server and can be done by using the following command:

```
# Chef-server-ctl install opscode-reporting
```

After installation, run the *reconfigure* command to configure and bring up the service. This will push the reporting data to the management service and enable its display in the Web UI. This is done by the following commands:

```
# Chef-server-ctl reconfigure
# opscode-reporting-ctl reconfigure
```

For each service, we need to run *Chef-server-ctl reconfigure* as well as the respective *service reconfigure*. This is to enable reconfiguration in both Chef Server and the

```
[root@magdesklinux ~]# chef-server-ctl test
Configuring logging...
Creating platform...
Configured URL: https://magdesklinux.wipro.com
Creating org pedant_testorg_magdesklinux_9929
Validating Org Creation
Run options: include {:focus=>true, :smoke=>true}

All examples were filtered out; ignoring {:focus=>true, :smoke=>true}

Randomized with seed 5034

.....................................

Finished in 2.16 seconds
33 examples, 0 failures

Randomized with seed 5034

Starting Pedant Run: 2015-09-16 13:53:17 UTC

 OPSCODE
 PEDANT
"Accuracy Over Tact"

--- Testing Environment ---
Config File: /var/opt/opscode/oc-chef-pedant/etc/pedant_config.rb
HTTP Traffic Log File: /var/log/opscode/oc-chef-pedant/http-traffic.log

Running tests from the following directories:
/opt/opscode/embedded/service/oc-chef-pedant/spec/api
Run options:
  include {:focus=>true, :smoke=>true}
  exclude {:intermittent_failure=>true, :cleanup=>true}

Randomized with seed 11990
```

Figure 1: Chef Server test

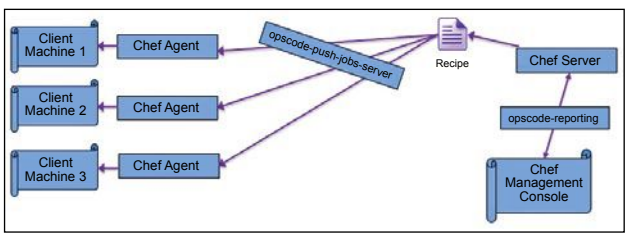

Figure 2: Chef job execution

respective service.

Preparing the client machine

As we are preparing the Chef set-up in a sandbox environment, we also need to prepare the client environments for the sandbox. While we are not setting up live environments with multiple nodes as agents with multiple machines, we need to set up virtual environments for client installation.

The sandbox environment is a single machine; the client VM can be another VMware image set up in the host machine (windows) or virtual machines on the guest VM (Linux). We have used an RHEL v7 based machine as the host to install the client VM.

Typically, a VM manager is part of the host system. If we didn't select the virtualisation packages while installing the server, we need to install *virt-manager* manually. To install the package, we can either use the RHELv7 ISO image, or set up the Yum repository by mounting the ISO file or by downloading it from the Red Hat website directly. This installation can be done by running the following command:

```
#yum install virt-manager
```

Client VM

Now, we can run *virt-manager* (as a sudo user) to install the client VM. Please follow the same steps we did to install the guest OS to prepare this client VM. Also note that we are taking the same guest OS resource—Chef Server and related services are running in the guest OS and, hence, use less resources for the client VM to enable enough resources to be available for the guest OS. The installation runs as shown in Figure 3.

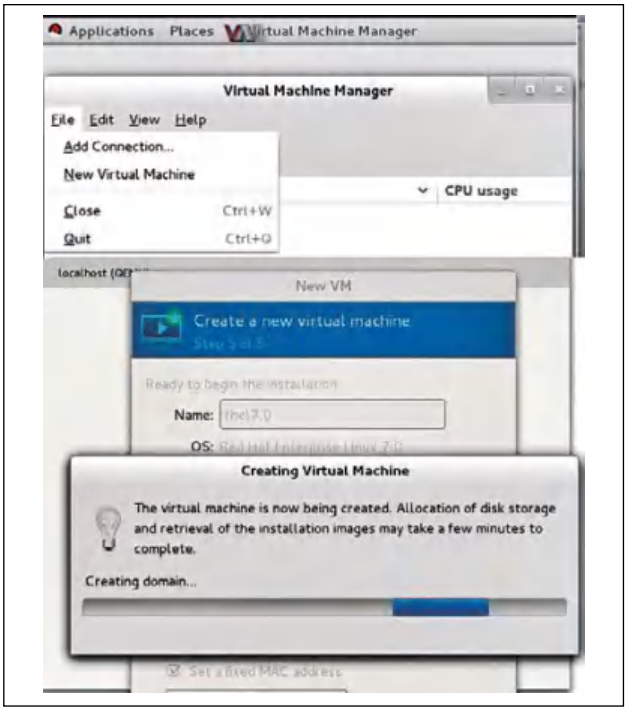

Figure 3: Installing the client VM

```
[root@magdesklinux ~]# virsh
Welcome to virsh, the virtualization interactive terminal.

Type:  'help' for help with commands
       'quit' to quit

virsh # edit Virt1
Domain Virt1 XML configuration edited.

virsh # start Virt1
Domain Virt1 started

virsh # list --all
 Id    Name                           State
----------------------------------------------------
 5     Virt1                          running

virsh #
```

Figure 4: *Virsh* command execution

Once the installation is over, edit */etc/libvirt/qemu.conf* and uncomment the user and group lines. This will enable the root user and root group permissions for *virt-manager*. If we do not do this, we may face permission issues while accessing the Linux image file when we try to run the client VM. Now start the virtual manager service using the following command:

```
# service libvirtd restart
```

Next, run the *virsh* command to verify if the client VM has

started successfully. This can be done as depicted in Figure 4.

We can also run the client VM from the *virt-manager* GUI. It is interactive and will guide you through the installation with the given ISO image. After selecting the ISO image, the installation procedure is the same as the installation of the host system as described in our previous article in this series.

Installing the Chef client

The Chef client needs to be installed in the client VM which we created using *virt-manager*. This process of preparing the Chef client is called bootstrapping. We can install Chef client by downloading it from *https://downloads.Chef.io/Chef-client/redhat/*.

Then, we need to install a tool called Ohai, which is used to synchronise various attributes to the nodes (client) at the start-up of the Chef client. Ohai collects data attributes from platform details, network usage, memory statistics, CPU usage, kernel health-check, host information including FQDN (fully qualified domain names) and statistical information. The private key created in Chef Server should be copied to all Chef clients so that two-way authentication (a handshake) happens between server (public key) and client (private key) when there is any communication between Chef Server and the client.

To download the certificate key from the server to all connected Chef clients, we need to run the *knife* command as follows:

```
#knife ssl fetch
```

This will download the certificate key from the server to the client in the folder */.Chef/trusted_certs*. To verify if key authentication between client and server has been successful, we can run the *knife* command in the Chef client node as follows:

```
#knife ssl check

Connecting to host magdesklinux.wipro.com:443 Successfully
verified certificates from 'magdesklinux.wipro.com'
```

Preparing *knife*

We have prepared Chef Server (on a Linux guest OS) and the Chef client (on a Linux client), so the next step is to prepare the communication channel between server and client. This is done with the configuration of *knife*, which is the command in the Chef environment to upload cookbooks, create clients, clone configuration between clients, upload and edit roles, and much more. *Knife* can be prepared by running the command as depicted in Figure 5.

As described in Figure 5, you should place the private key (*pem* file) in the */etc/Chef-server* folder. Once this configuration is done, we can verify that the handshake of keys is happening and the client is configured properly by executing the following command:

```
#knife client list
```

```
[root@magdesklinux ~]# knife configure
WARNING: No knife configuration file found
Where should I put the config file? [/root/.chef/knife.rb]
Please enter the chef server URL: [https://magdesklinux.wipro.com:443]
Please enter an existing username or clientname for the API: [magesh]
Please enter the validation clientname: [chef-validator]
Please enter the location of the validation key: [/etc/chef-server/chef-validator.pem]
Please enter the path to a chef repository (or leave blank):
*****

You must place your client key in:
  /root/.chef/magesh.pem
Before running commands with Knife!

*****

You must place your validation key in:
  /etc/chef-server/chef-validator.pem
Before generating instance data with Knife!

*****
Configuration file written to /root/.chef/knife.rb
[root@magdesklinux ~]#
```

Figure 5: *Knife* command

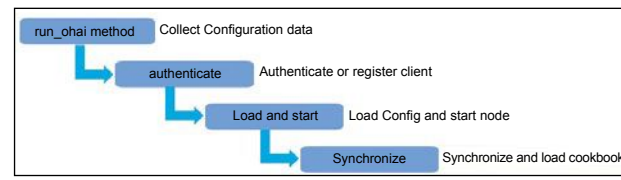

Figure 6: Chef client start-up steps

Now, when we bring up the Chef client, the following basic steps (also shown in Figure 6) happen in the background during client start-up. Please note that these are major steps in the client start-up, though there are many smaller steps involved in between these steps. This stage is explained in Figure 6.

run_ohai: This is the first step when we start the Chef client, where it loads the configuration data as listed earlier. This is collected for each node of the Chef environment and starts the Chef client.

Authenticate/Register: This step is most important to bring up the Chef client into the secured environment of the Chef set-up. If the authentication is successful through the two-way handshake (public/private key authentication), the client node is registered to the Chef Server so that it is entered into the Chef environment and is available in the reporting UI of the Chef management console.

Load and run node: Here, the configuration is loaded, and the client is authenticated and registered to the Chef Server environment. The next step is to load the node and start the node machine.

Synchronizecookbooks: This synchronises the cookbooks from the Chef Server. The cookbooks are stored in the storage of the Chef Server and loaded into the client machines through the *Knife* tool (pushed to all client nodes).

After this, the set-up is finalised. This saves the node state and finalises the client set-up. If any of these steps fail, the Chef client moves to *start-up run_failed* state and stops abnormally. We need to look into the log file to analyse the reason for the failure and fix it.

Now, run *runs list* to get the list of RUN_ID of the Chef client which is successfully running. This is done by executing the following command:

```
# knife runs list
```

Continued on page 56...

The Recommender System for Big Data

In social networking sites and in e-commerce, recommendations play a key role. E-commerce, in particular, depends on website traffic which is driven to it by various means. Fortunately, machine learning algorithms like Recommender can be used effectively in both these cases. This article focuses on Mahout's recommender systems.

Mahout is an open source machine learning library from the Apache Software Foundation. It implements many data mining algorithms like *recommender engines(), clustering (), classification ()*, and is scalable to very large data sets—up to terabytes and petabytes, which is in the Big Data realm.

In this article, I will focus on recommender systems in Mahout. Recommender engines are very popular machine learning algorithms that are used to recommend books, movies or articles based on the users' past actions and interests.

Recommender engine algorithms come in two categories – user based and item based recommendations. User based recommendations predict what the users will like, based on their similarity with other users and do not require the properties or description of the items; for example, Facebook's friend recommendations. In item based recommendations, the preferences of users for some items based on their preference of other similar items can be suggested, and this requires prior knowledge of the particular item's properties. For example, if a user likes one kind of comedy movie, he may like other similar comedy movies.

The following are the main components of the recommendation system in Mahout.

- *Data model interface:* This converts the raw data to a Mahout compliant data format. Data sources like MySQL, PostgreSQL, MongoDB, Cassandra, flat files, etc, are supported.
- *User similarity interface:* This contains methods to compute the similarity among users.
- *Item similarity interface:* This contains methods to compute the similarity between items. Many similarity metrics are available in Mahout like Pearson Correlation, Euclidean distance, etc.
- *User neighbourhood interface:* This is a method to construct a neighbourhood around a given user, which satisfies the similarity or 'nearest neighbour' threshold criteria.
- *Recommender interface:* This is a method that finally makes the actual recommendation.

In order to process Big Data, Mahout can use HDFS. So a prerequisite to installing Mahout is JDK/JAVA 1.7 or later, Maven 3.0 or later and the Hadoop cluster. If you are using Ubuntu, you can use the following commands.

Install Maven from the repository *sudo apt-get install maven.* Download the latest distribution of Mahout from the site *http://www.apache.org/dyn/closer.cgi/lucene/mahout/*. Unzip and copy the following command to the desired location:

```
cp -R SOURCE  DESTINATION
```

I copied it in */usr/local*.

Now, cd */usr/local/mahout/distribution* or wherever you have copied Mahout. Then run the following command:

```
sudo mvn install
```

(Install Maven 3.0.1 or above for Mahout .20 distribution; else, it will throw some error.) You will get the screen shown in Figure 1 after successful installation.

Mahout's recommenders expect interactions between users and items as input. I tested a sample 568.3MB data,

which contained the fields *userID*, *movieID* and *value*. Here, *userID* and *movie ID (itemID)* refer to a particular user and a particular item, and *value* denotes the strength of the interaction (e.g., the rating given to a movie).

The following steps can be used to run the Recommendation algorithm.

Create a directory in the Hadoop file system to store the ratings file using the following command:

```
hadoop fs -mkdir /mahout_data/
```

Copy the downloaded file to HDFS using the following command:

```
hadoop fs -put /home/hduser/mydata/ml-latest/
ratings.csv /mahout_data/
```

Go to the Mahout directory, cd */usr/local/mahout/bin/* and issue the following command (the output file should be unique and JAVA_HOME should be properly set):

```
./mahout recommenditembased -s
SIMILARITY_LOGLIKELIHOOD -i hdfs://
localhost:9000/mahout_data/ratings.csv
-o hdfs://localhost:9000/ratings_test/
--numRecommendations 25
-i hdfs://localhost:9000/mahout_data/ratings.csv - Denotes
the input file
-o hdfs://localhost:9000/ratings_test/  -denotes the
output file .
```

The command will run for a couple of minutes and you can see your output from the Web interface as well, as shown in Figure 2.

Note: In the above snippet, *recommenditembased* means we are creating an item based recommendation and not a user based recommendation. The difference between the two is that a user based recommendation finds similar users based on what they like, and item based recommendation figures out what the user likes and finds items to match those preferences. Mahout's item-based recommendation algorithm takes customer preferences by item as input, and generates an output recommending similar items with a score indicating whether a customer will *like* the recommended item.

You can check the output file and it will contain two columns: the *userID* and an array of *itemIDs* and scores. This could, for instance, recommend a user's preference for a particular movie, which he may be interested in watching.

Figure 1: Building Apache Mahout on Ubuntu

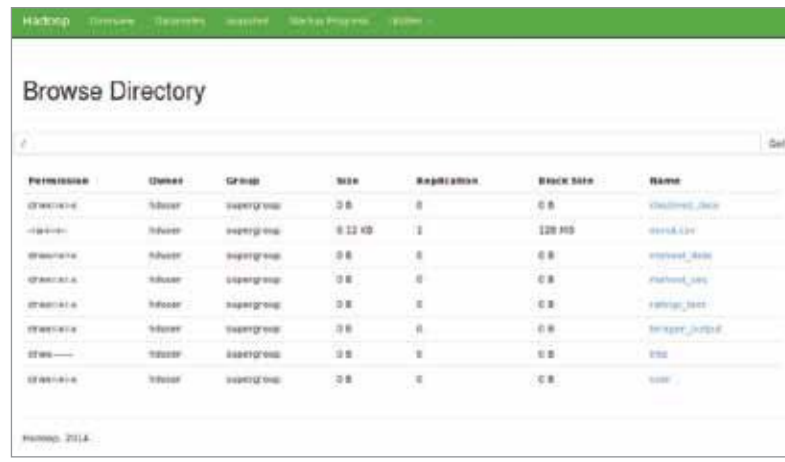

Figure 2: Browsing directory

"We are leaving the age of information and entering the age of recommendation," says Chris Anderson, British entrepreneur and the curator of TED. The increasing adoption of the Web as a vehicle for business has changed the way in which businesses interact with the customer. Marketing teams and businesses are now using intelligent algorithms and social technologies to form meaningful, ongoing relationships with customers. It is here that technologies like Mahout Recommender will play a key role. END

References

[1] *http://mahout.apache.org/users/recommender/intro-itembased-hadoop.html*
[2] *http://grouplens.org/datasets/movielens/*
[3] *http://info.mapr.com/rs/mapr/images/PracticalMachineLearning.pdf*
[4] 'Mahout in Action' by Sean Owen, Robin Anil, Ted Dunning, Ellen Friedman (Manning Publications Co.)

By: Surabhi Dwivedi

Surabhi Dwivedi has an experience of nine years in the areas of software development, database technologies, data modelling, data mining, Linux systems administration, open source technologies and Big Data.

Configuring an SDN Controller in Open Source Mininet Emulator

The exponential increase in the use of the Internet, coupled with the growth in cloud computing, has created a need for better network control and management. Software defined networking (SDN) – a means to control the network through software – is a solution to the problem. This tutorial, aimed at researchers working in the SDN and the cloud computing domain, presumes knowledge of Mininet Emulator, Python and SDN and would be of interest to B. Tech, M. Tech and PhD scholars looking for an open source simulation platform for SDN.

SDN is a revolutionary idea in computer networking that ensures significant flexibility and simplicity in network control and management, apart from giving a broader scope for innovation through programmability. SDN is all about controlling the network through software or, in other words, making the network programmable.

There has been a dramatic increase in Internet usage over the past few years, as compared to the use of networking technology, which hasn't been engaged as much. SDN has been introduced as a replacement for conventional networking, to meet market requirements. Fortunately, SDN couldn't have evolved at a better time. Since the use of cloud computing is increasing, it is necessary to automate the configurations as much as possible.

According to the ONF (Open Network Foundation) definition, SDN is actually a decoupling of the network control plane and the forwarding plane. With the demand for cloud computing increasing, SDN has evolved as the most efficient way of controlling networking using some high level language to make the programming as flexible as possible.

Traditional networking versus SDN

The main difference between traditional networking and SDN is the way in which data is handled and forwarded. Unlike traditional networking, SDN has separate devices called SDN controllers which control the data path. These control the path of packets coming to switches known as the OFvSwitch (Open Flow Virtual Switch). So, there has to be some interface between the forwarding plane and the controlling plane, which is provided by the OF (Open Flow) protocol. SDN enables admins to control the way switches handle the data, provide QoS (Quality of Service), and automate the process to make it less tedious and erroneous.

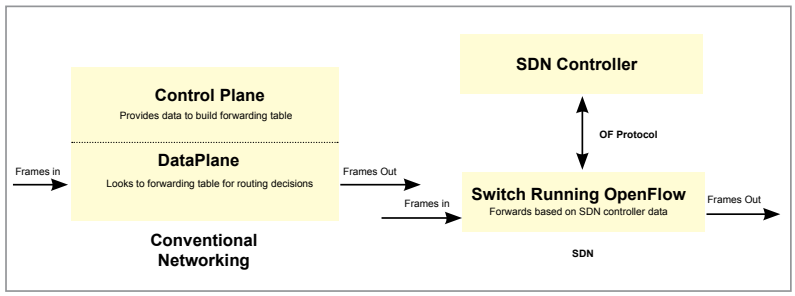

Figure 1: Abstract view of conventional networking versus SDN

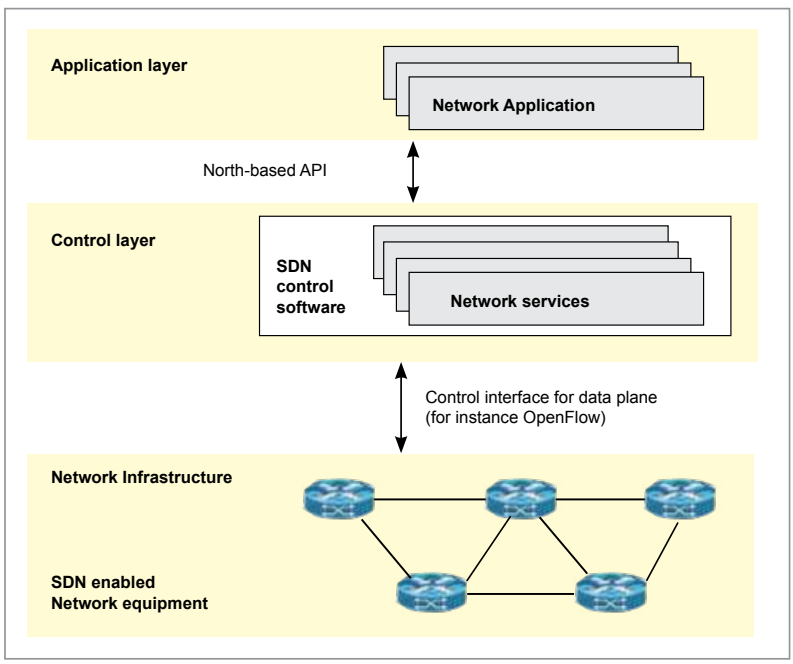

Figure 2: SDN or software defined networking

Advantages of SDN over traditional networking

SDN has many advantages over traditional networking:

a. Due to the introduction of some automation in the process of networking through SDN, scalability has been increased significantly, which is also a critical requirement of the current market.
b. Unlike conventional networking, SDN only requires one centralised control plane which offsets the cost of the forwarding plane.
c. VM migration becomes easier.
d. Automating the configuration is possible.
e. Quality of Service can be provided in a more efficient way.

As shown in Figure 2, the SDN controller provides a programmable interface to the OF switches. With the help of this interface, different network applications can be written to control, manage and offer new functionalities. A recent study of several OpenFlow implementations, conducted on a large emulated network with 100,000 hosts and 256 switches, revealed that all the controllers were able to handle about 50,000 new flow requests per second. Besides, according to research going on at Stanford University, the new architecture will support around 20 million requests per second with about 5,000 switches. Multiple controllers can be used for scalability purposes; these also allow backup controllers to overcome failures and provide almost 100 per cent uptime.

Configuring an existing basic controller (POX controller) in Python

POX is an open source development platform for Python-based software defined networking (SDN) control applications, such as OpenFlow SDN controllers. POX, which enables rapid development and prototyping, is now being more commonly used than NOX, a sister project which is in C/C++.

In this article, we'll configure a basic SDN controller (POX) in Python, and emulate using an open source emulator called Mininet, which provides the functionality to create and operate a virtual network and control it using a controller. To set up the environment for a Mininet emulator, follow the steps described in the article that appeared in the September 2015 issue of *OSFY*. Quick commands to run on the terminal are given below for installation purposes (these commands have been tested on Ubuntu 14.04 and might change for other versions or distributions):

```
sudo apt-get install mininet
```

For installation from the source, type:

```
git clone git://github.com/mininet/mininet
cd mininet
mininet/util/install.sh [options]
```

…where options include the installation of various packages along with Mininet, which can be seen using the *Help* option.

After the successful installation of Mininet (along with the POX controller), we can start POX using the following command:

```
./pox.py log.level –DEBUG misc.of_tutorial
```

The above command tells the POX controller to enable verbose logging and to start the *of_tutorial* file, which acts as an Ethernet hub right now.

Now, start the *Mininet openflow* tool to perform experiments using the following command line:

```
sudo mn –topo single, 3 --mac --switch ovsk --controller
```

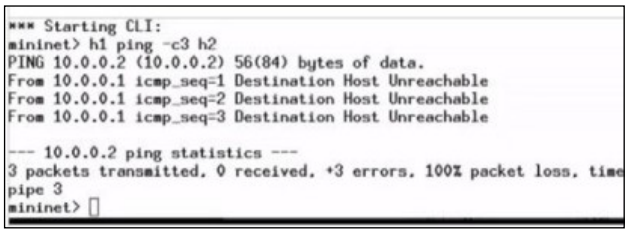

Figure 3: Remote controller without any entries in switch and thus dropping all the packets

```
remote
```

After successfully starting, it will show the message that switches have been connected with the *{MAC address}*. To verify the connections established by default, use the following command:

```
pingall
```

Here, all the hosts will be unreachable to one another. This is because we have set the remote controller, but have not initiated it or, in other words, there is no entry in the controller for how to handle incoming packets at the switches. Thus, when OF-enabled switches ask for the decision from the controller, it simply gives the command to drop the packets. Figure 3 shows a snapshot of the following command:

```
h1 ping -c3 h2
```

We can now control each and every host of the network in the emulator, virtually. For that, we can use the following command:

```
xterm h1 h2 h3
```

Three small terminal-like windows will pop up. To check that everything is working fine, we can just capture the packages. Assuming that support for Mininet in Wireshark is also installed, run the following commands:

```
tcpdump -XX -n -I h1-eth0
tcpdump -XX -n -I h2-eth0
```

Now, we are monitoring the traffic on *h1* and *h2* hosts. Let's run the ping command in *xterm* of *h3*.

```
ping -c1 10.0.0.1
```

After running this command, we can see the ARP requests in the *h1* terminal.

Customising an existing POX controller to act as a firewall and load balancer

Till now, we were making the controller work like a hub. To make it work like a switch or learning switch, we can make changes in the file *of_tutorial*, and replace the statement *act_like_hub()* to *act_like_switch()*.

An SDN based firewall

For a firewall kind of application, we need to have a mediator that can filter out the packets based on some conditions. Let us consider an example with MAC address filtration. For that we require to build a look-up table that can hold the MAC addresses of all the devices allowed to communicate or transfer data.

To make the controller communicate with switches, the former needs to send the messages to the latter. As and when a connection initiates a switch, even *ConnectionUp* is fired. Thus, tutorial code creates a *Connection* object, which can later be used to send the messages to the switches using the function *connection.send()*. At that time, we can decide the default flow using the *ofp_action_output* function, as shown below:

```
out_action = of.ofp_action_output (port = of.OFPP_FLOOD)
```

Here, *OFPP_FLOOD* even decides on flooding, which means that the packets will be forwarded to every port except the one on which the packet had arrived.

A firewall needs to restrict the packets on the basis of several other criteria. For that, we need to create the object *ofp_match* class. Some important fields of this class are *dl_src*, *dl_dst* and *in_port*. Here, *dl* represents the data link layer, which means *src* and *dst* are MAC addresses of the source and the destination, respectively. To send the packet, we need to send a message using *ofp_packet_out*. The *send_packet()* method can be used in *of_tutorial*.

Here is an example:

```
def send_packet (self, buffer_id, raw_data, out_port, in_
port):
        """Sends the packet out of the specified port.
```

If *buffer_id* is a valid buffer on the switch, use it; otherwise, send the raw data in *raw_data*.

The *in_port* is the port number that the packet arrived on. Use *OFPP_NONE* if the packet is generated by you.

```
        """
        msg = of.ofp_packet_out()
        msg.in_port = in_port
        if buffer_id != -1 and buffer_id is not None:
                msg.buffer_id = buffer_id
        else:
                if raw_data is None:
                        return
                msg.raw_data = raw_data
    action = of.ofp_action_output (port = out_port)
        msg.actions.append (action)
        self.connection.send(msg)
```

This function gives the command to the switch to enter the flow table entry. Now, to take appropriate action at the switch, we need to create the entry for the packet that we want to route. Let's say we want to route a packet coming to input port = *2*. So, we need to create a matching object for that, which is given by *ofp_match*. To create the entry in the switch, write the following code:

```
fm = of.ofp_flow_mod()
            fm.match_in_port = 2
            fm.actions.append(of.ofp_action_output(port = 4))
```

Thus, as and when a packet comes at the port = *2*, it will be redirected to port = *4*. We can definitely add more parameters as per the requirement, like *idle_timeout, hard_timeout, actions, priority, buffer_id, match, in_port*. But we have not included all of them -- just to avoid complexity.

An SDN based load balancer using round robin scheduling

As writing a load balancer could be a tedious task, we can use a template of an existing controller which forwards a request to the available servers randomly. The template can be found at *https://github.com/noxrepo/pox/blob/carp/pox/misc/ip_loadbalancer.py*

The main role of the controller is to select the server to which the incoming requests are to be forwarded. This part is coded in the function of the template called *pick_server*. It is found to be at Line No. 190 in the standard template. Thus, we can modify that procedure to customise the default policy.

The code below refers to Line No. 190 of the file *ip_loadbalancer.py:*

```
def _pick_server (self, key, inport):
"""
Pick a server for a (hopefully) new connection
"""
return random.choice(self.live_servers.keys())
```

The above lines show the random selection of the servers for forwarding. Now we are going to change this method to remove the randomness and add the round robin algorithm to the load balancer. First, we will add a new variable called *selected_server* to class *iplb* (which is the class containing the above method, i.e., *pick_server()*). The *selected_server* will be an instance variable defaulting to 0, which will keep track of the state of the server allocated to fulfill the immediate previous request.

So, a modified class definition of *iplb* will be as shown below:

```
class iplb:
              def __init__(self,...):
...
self.selected_server = 0
```

Here, *__init__* is the constructor and the *self* inside the method refers to the instance of the object. So is the case with anything defined with *self. Prefixed* can be used as an instance variable (similar to other object-oriented languages such as Java). So now, to change the *pick_server* method, run the following commands:

```
def _pick_server (self, key, inport):
"""
```

Pick a server for a (hopefully) new round robin based connection.

```
"""
all_keys=self.live_servers.keys()
 if self.selected_server==len(self.live_servers):
self.selected_server=0
redirected_s=all_keys[self.selected_server]
self.selected_server+=1
return redirected_s
```

The above code first gets all the keys from the *self. live_servers* (which is a Python dictionary, i.e., a Hash Map). Then, in the next line, it makes sure that the *selected_server* isn't out of range to the total number of servers. Finally, it gets the key of the selected server and then increments the variable for the next call. END

References

[1] *http://www.blackstratus.com/blog/traditional-software-defined-networking/*
[2] *https://globalconfig.net/software-defined-networking-vs-traditional/*
[3] *https://www.opennetworking.org/sdn-resources/sdn-definition*
[4] *https://openflow.stanford.edu/display/ONL/POX+Wiki*
[5] *https://github.com/noxrepo/pox*
[6] *https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf*
[7] *http://mininet.org/sample-workflow/*

By: Jitendra Bhatia, Shreyans Doshi and Girish Ramnani

Bhatia works as a senior assistant professor with the CSE department of the Institute of Technology, Nirma University for the past nine years. You can contact him at *jitendra.bhatia@nirmauni.ac.in*

Doshi is a B. Tech final year student in the CSE department at the Institute of Technology, Nirma University. He can be contacted at *Shreyans.doshi94@gmail.com*

Ramnani is a B. Tech pre-final year student in the CSE department at the Institute of Technology, Nirma University. You can contact him at *girishramnani95@gmail.com*

Gift 'Knowledge' To Your Friends and Your Near and Dear!

On birthdays, anniversaries and occasions like Diwali, New Year or Eid, you may be looking for gifts that are useful for the recipients and also last long. If so, what could be better than gifting a magazine that they may like and that would add to their knowledge, month after month.

If the person you want to gift to is an electronics engineer or likes the subject as an industrialist, academician, student or a hobbyist, you may consider gifting him Electronics For You, the most popular magazine on the subject. He/she will love you for it – every time he/she looks at it.

If the person you have in mind is interested in electronics as a businessman or as a non-technical person, gift him/her subscription to Electronics Bazaar, which is far above the rest in terms of the latest and credible information it offers.

For those interested in software, what can be better than Open Source For You magazine. It would keep the person updated with the latest in the field of 'open' and free software available – saving a lot of money in the long run. It comes with a DVD packed with very useful information/software every month.

What's amazing is that gift subscription to any of these magazines would cost you much lesser then what you would normally be paying for most other gifts! Please check the amount below that you have to pay for 12 issues, including packaging and forwarding:

Magazine	You Pay (Rs)
Electronics For You*	☐ 960
Electronics Bazaar**	☐ 960
Open Source For You	☐ 960

*For Electronics For You Plus edition that includes a very useful DVD every month besides delivery by courier

**It includes LED Bazaar

The subscription can start from any issue/month you desire, though normally a magazine is mailed from the next issue being sent to all the readers.

One more thing: Access to digital version of the magazine(s) will also be provided to the person you are gifting to, so that they do not miss the information in case of any postal delay in receipt of their copies.

And, if you are not subscribing to the magazine(s) yourself yet, don't forget to gift it to yourself first!

Please note, these rates are likely to increase soon! So please send your order before the increase to avail these rates.

To Subscribe Online, Visit: http://subscribe.efyindia.com

DONOR:

Name_____ Mailing Address_____

_____ City_____ Pin Code_____

State_____ Phone_____ Email_____

I would like to gift subscription to the above (✓) marked magazine(s) starting from_____. Please find enclosed a sum of Rs_____ by DD/MO/crossed cheque*bearing the No._____ dt._____ in favour of **EFY Enterprises Pvt Ltd**, payable at Delhi. (*Please add Rs 50 on non-metro cheque)

RECIPIENT MAILING DETAILS:

Name_____

Mailing Address_____

_____

City_____

Pin Code_____ State_____

Phone_____ Email_____

EFYGROUP Technology Driven Us

EFY Enterprises Pvt Ltd
D-87/1 Okhla Industrial Area, Phase 1, New Delhi 110020

Phone:
011-26810601 – 03

Email:
info@efy.in

www.efy.in

Exploring Processes Using Sysinternals

Sysinternals is a freeware tool that can help to manage, diagnose, troubleshoot and monitor a Windows environment. It is a part of the Microsoft TechNet website which offers diagnostic tools, technical resources and utilities. This article briefly explains how Sysinternals can be used to give information about running processes.

You can now use the Sysinternals tools to find out which program has a particular file or directory open. The tools available in the Sysinternals suite show details about running process handles and DLLs. The entire suite can be downloaded from *https://download.sysinternals.com/files/SysinternalsSuite.zip*

Process monitor

Process monitor is a free utility available to monitor real-time file systems, registry, process/threads and the network. It can be downloaded from *https://technet.microsoft.com/en-us/library/bb896645.aspx.*

Trace summary tools

Process monitor includes a number of dialogues that allow you to perform simple data mining on the events collected in a trace.

Process tree displays all of the processes referenced in a hierarchy in the loaded trace, which shows parent-child relationships. Go to *Tools > Process Tree* or *Ctrl+T.*

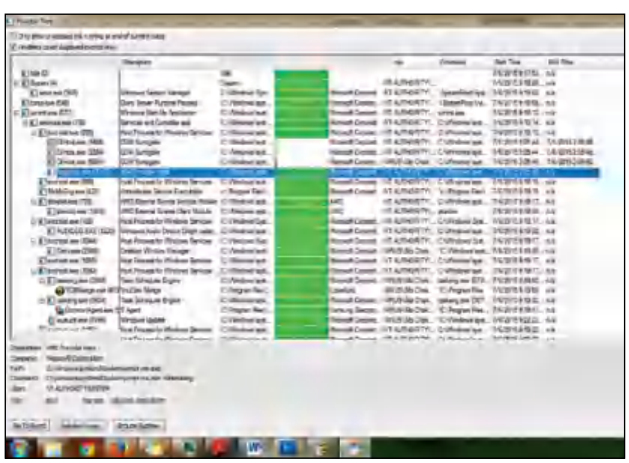

Figure 1: Process tree for various processes

File system

Process monitor displays all the activities of a file system, including local and remote storage. It also detects and monitors new file system devices.

File summary

File summary dialogue lists each unique file system path present in the filtered trace, the amount of time spent performing I/O to the file, the total number of events that referenced the path, and the count of individual operation types.

File summary can be accessed by going to *Tools > File summary.*

The file summary can be achieved by means of the folder and extension.

Activity summary

Activity summary lists all the processes seen in the trace, file events, I/O, registry events, network events, including their process ID, image name and command line.

Activity summary can be accessed by going to *Tools > Process Activity Summary.*

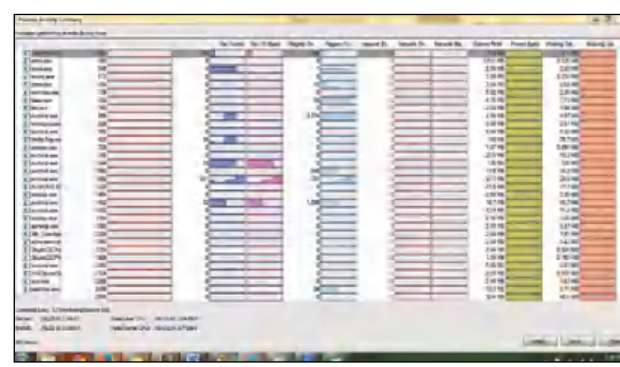

Figure 4: Activity summary for various processes and their operations

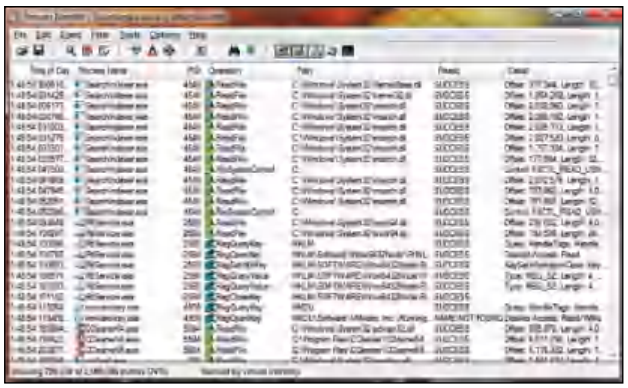

Figure 2: File system activity

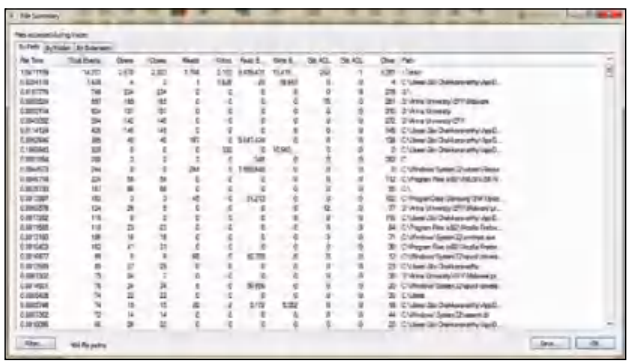

Figure 3: File summary by path

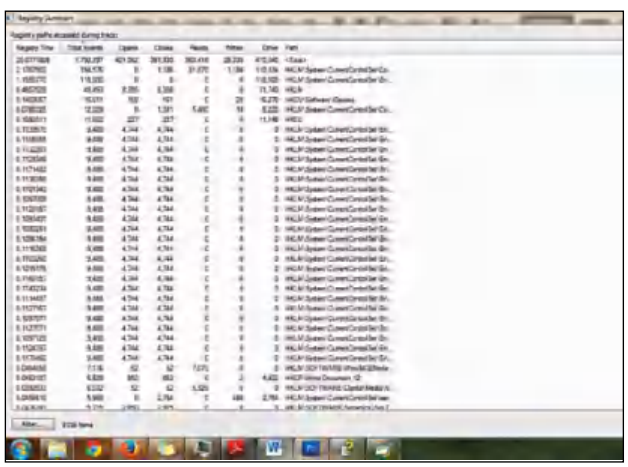

Figure 5: Registry information accessed during trace

Registry summary

Registry summary lists each unique registry path present in the filtered trace, the amount of time spent performing I/O to the registry path, the total number of events that referenced the path, and the count of individual operation types.

Registry summary can be accessed by going to *Tools* > *Registry Summary.*

Stack summary

Stack summary is used to visualise individual instances of stack traces for each process. You can access the stack

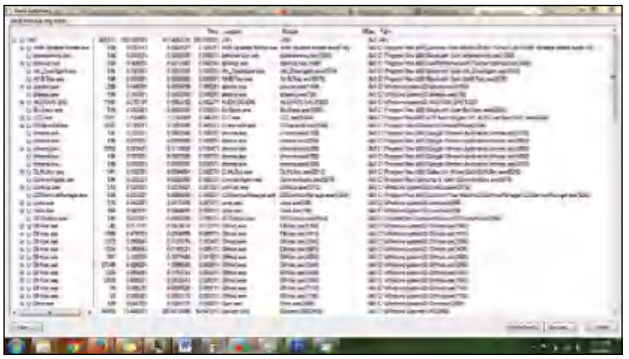

Figure 6: Stack information during trace

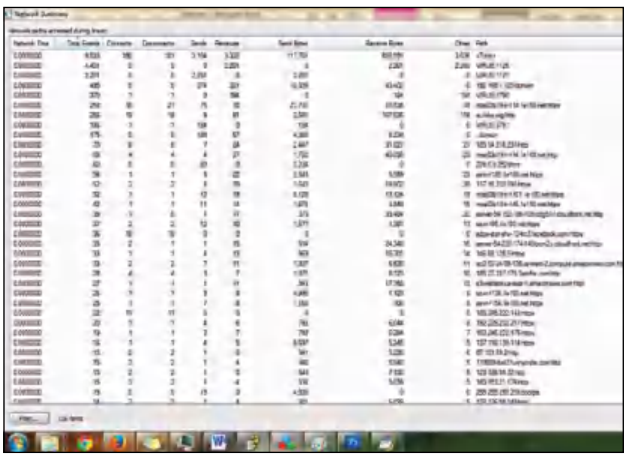

Figure 7: Network information during trace

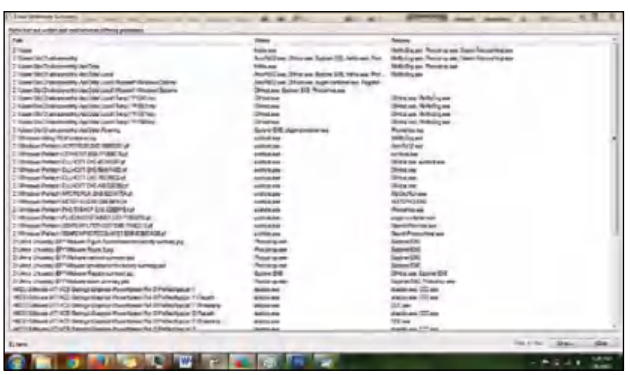

Figure 8: Cross reference summary during trace

summary by going to *Tools* > *Stack Summary.*

Network summary

Network summary lists each unique destination IP address present in the trace and a number of different types of events, including sends and receives, to each address.

Go to *Tools* > *Network Summary.*

Cross reference summary

This dialogue box shows the paths that are written by one process and read by another one.

Go to *Tools* > *Cross reference summary* (paths that are

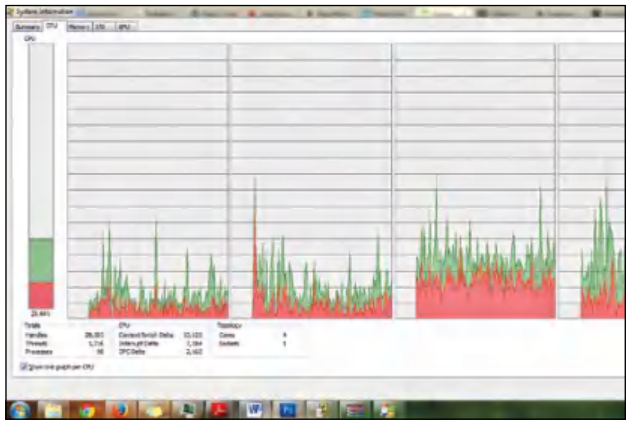

Figure 9: Visualising the CPU process

written and read between differing processes).

Visualising the process using Process Explorer

Process Explorer helps to visualise the process which, in turn, helps to deeply observe the process and its handles. It can be downloaded from *https://technet.microsoft.com/en-in/sysinternals/bb896653.aspx*

The tools explained above can help you to learn and analyse the behaviour of a process/thread. The process has to be carried out manually. END

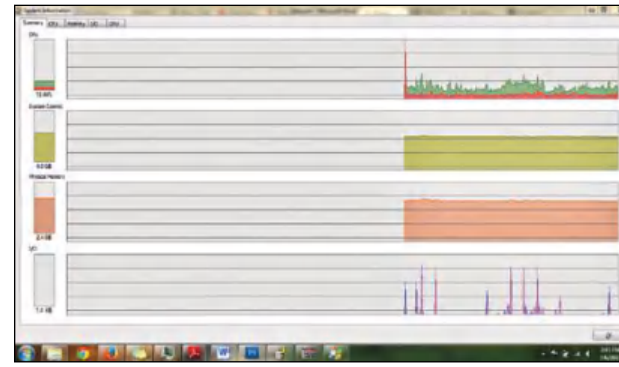

Figure 10: A visual summary of all the processes

References

[1] *https://technet.microsoft.com/enin/sysinternals/bb896653.aspx*
[2] *https://download.sysinternals.com/files/SysinternalsSuite.zip*

By: Sibi Chakkaravarthy

The author holds an M. Tech degree in computer science and engineering and is working on a PhD at Anna University. Currently, he is with the AU-KBC Research Centre, MIT, Chennai. His research interests include network security, wireless security, wireless sensor networks, cloud security, etc. He can be reached at *sb(DOT)sibi(AT)gmail(DOT)com.*

Continued from page 46...

From the output of the above command, you can get the RUN_ID of the client and run the *show* command to get client details. For example:

```
# knife runs show c40csacd2-6778-122d-45589-3e1c4e55f741
```

Execution of the above command will display the running information of the client like node name, ID, status, run resources, checksum, permissions, user and other relevant details. With this step, the client is successfully installed, configured and able to communicate with the registered server.

Preparing the reporting facility

Now the client is completely ready, and we need to enable its reporting facility so that reporting data is collected in the server and displayed in the management console. This can be done by editing the *client.rb* file and setting the option *enable_reporting true*. At the same time, run the reporting service in Chef Server (Guest OS) as follows:

```
#opscode-reporting-ctl start
```

To get some sample reporting scripts written in Ruby, you can download them from the Git repository *https://github.com/Chef-cookbooks/Chef-reporting.*

Use the *git fetch* command to download these Ruby scripts and add *Chef-reporting::default* to *client.rb* to enable the publishing of reporting data. There's no problem in introducing Ruby scripts suddenly.

So far, we have looked at how to set up a client VM to install the Chef client, and how to install a Chef client and Chef reporting service. Also, we have learnt the purpose of the *knife* tool, as well as how to install and configure it. In our next article, we will discuss Ruby scripts and their implications in Chef Cookbook preparation. Also, we will cook up some recipes for Chef configuration, and learn how to do a typical deployment of the application and how Knife is used in the Chef environment.

In subsequent articles in this series, we want to showcase a typical real-time banking/financial infrastructure to show how Chef helps in automating the deployment process with just a single button click, as well as how it helps in resilience and automating scaling of infrastructure to meet customer demands. END

By: Magesh Kasthuri and Dr B. Thangaraju

Kasthuri is a senior distinguished member of the technical staff at Wipro. He is a senior technical consultant in the Java-Technology Practices Group and can be reached at *magesh.kasthuri@wipro.com.*

Thangaraju is an open source software evangelist, who works at Talent Transformation, Wipro Technologies, Bangalore and can be reached at *balat.raju@wipro.com.*

JavaScript: The New Parts

This article is the fourth part in this series. The author discusses the *for...of* iterator, collections and arrow functions.

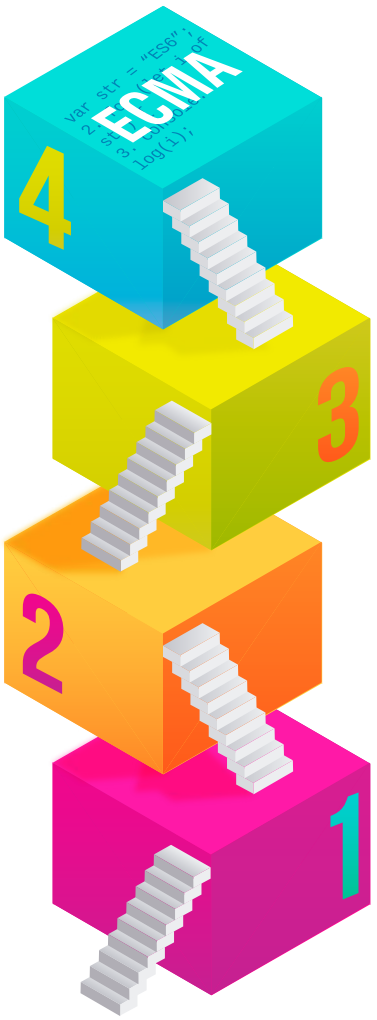

The JavaScript language is evolving to enable the development of complex and large applications. This is a multi-year effort by ECMA (a standards body). In the previous article or third part of this series, we learnt about computed property keys and string functions. Their support is available in most recent versions of browsers. In this part, we will learn about a control structure for collections, new collection data structures and abbreviated syntax for function-definition using the arrow symbol (=>). The arrow function is an exciting feature of ES6, loved by most JS enthusiasts.

The *for..of* iterator

JavaScript has *for* and *for...in* control statements to loop through a block of code a number of times and traverse through properties of an object, respectively. Note that *for...in* loops through the properties of an object. In ES6, a new looping structure *for...of* is introduced to loop through values of a collection. In addition to strings and arrays, we have Maps and Sets introduced in ES6.

Snippet 1: Example of an existing for..in *and new* for..of *looping statement*

```
1. let arr = [“mon”, “tue”, “wed”];
2. for (let i in arr) {
3. console.log(i); //prints index
4. }
5. for (let i of arr) {
6. console.log(i); //prints values
7. }
8. for (let [i,v] of arr.entries()) {
9. console.log(i, “: “, v); // prints index, values
10.}
```

Output:

```
0
1
2
mon
tue
wed
0: mon
1: tue
2: wed
```

In Lines 2-4, we have the looping structure, which displays only the index in case of arrays and property names in case of objects. In Lines 5-8, the new *for..of* statement assigns the value of the array to *i*. Lines 8-10 show another variation, where we have both the index and the value captured.

Snippet 2: String object iterated using for...of *loop*

```
1. var str = “ES6”;
2. for (let i of str) {
3. console.log(i);
4. }
```

Output:

```
E
S
6
```

for..of can also be used to iterate through the arguments object of a function/method. Note that *arguments* is a local variable available within all functions.

Snippet 3: Example to demonstrate arguments object

```
1. function fn(one, two, three) {
2.for(let arg of arguments) {
  console.log(arg);
3.}
4. }
5. fn("mon","tue","wed");
Output:
mon
tue
wed
```

```
let makept = (a, b) => ({ x:a, y:b});
```

Arrow function that returns an object.

```
let capcat = (a, b) => {
 let c = a.toUpperCase()
 c += " " + b.toUpperCase();
 return c;
}
```

Arrow function with body and an explicit return at the end.

In addition to strings, arrays, arguments and new collections, *for...of* can also be useful with DOM manipulation.

To provide this capability of iterating through values, built-in functions have been added to the array object, i.e., the prototype method of string and array objects has the *Symbol.iterator* method. This new *Symbol.iterator* introduction, in combination with other new features such as generators, makes this new *for...of* statement powerful.

Collections

Maps and Sets are two new categories of collections added to ES6. Maps allow the storing of arbitrary values as a collection, and Sets is used for storing a collection of unique elements.

Map is composed of key-value pairs. Language supports built-in operations of *set()*, *get()*, *has()* and *delete()*.

Set has operations of *add()*, *clear()* and *has()* to add elements, delete all elements and check the presence of an element, respectively.

Two other variants of Map and Set are the WeakMap and WeakSet. The Weak variants of collections have a few restrictions. Weak variants cannot be iterable, which means we cannot use the *for...of* statement. The second limitation is that we can only store objects and not primitive data types.

One reason for introducing WeakMap and WeakSet is automatic garbage collection. In the case of Map and Set, the programmer needs to take responsibility for checking references and clearing the memory.

Arrow functions

A new shorthand notation for anonymous functions is the arrow function. It is represented by the symbol =>. This notation of equal-to instead of a dash (or hyphen) makes it a fat arrow function. This notation has been borrowed from.CoffeeScript, and is extremely useful when defining callback functions. Arrow functions have more than cosmetic features. They overcome some of the side effects of *this* binding with traditional functions. In addition, they also help JS engines to perform better with something called 'tail call optimisations'. For details, go to *References* at the end of the article.

Let's look at an example of an anonymous function definition like the one below:

```
1. var std = function std(yr) {
2. return "ECMAScript" + yr;
3. };
4. console.log(std("2015"));
```

With the ES6 arrow function, the same can be shortened to:

```
1. var std = yr => "ECMAScript"+yr;
2. console.log(std("2015"));
```

Output:

```
ECMAScript2015
```

Points to note in the new syntax are that the keyword function has been replaced by the => symbol, and there are no parentheses around the parameter. Since the method has only one statement (expression), there are no curly brackets to enclose the function body. There is no return statement as the body of the function is an expression.

Given below are a few examples of arrow functions, with varying number of parameters and body of the function.

```
let x = () => {};
```

Arrow function with no parameters and no body.

```
let y = (a, b) => a + b;
```

Arrow function with parameters in parenthesis and function body is an expression.

One major advantage of using arrow functions is the binding of *this*. It is well known that in JavaScript, there is the annoyance of the *this* keyword taking different values based on invocation. In contrast, in arrow functions, *this* has lexical scoping, which means there is no binding. The value of *this* is determined by the containing function. This solves many side effects in callback handlers. Developers who have used libraries like jQuery should be more familiar with this issue.

With these advantages, it might seem that arrow functions are clear replacements to regular functions. While this makes sense in most cases, some differences remain.

Comparison with traditional functions

In comparison with traditional functions, arrow functions cannot be used as constructors. The value of *this* is fixed and cannot be changed. The local object *arguments* is not associated with the arrow function, but the calling function.

List processing

```
[1,2,3,4].map(function (x) {
 return console.log(x * x);
});
```

In ES6

```
[1,2,3,4].map(x => console.log(x*x));
```

Features support

The features covered this month are supported by the latest version of Chrome, Firefox and Node 4.0 environments. For more granular feature support, refer to Kangax matrix in *References*.

In the next part in this series, we will cover the set of features related to function parameters. These are: default function parameters, rest parameters and the spread operator. END

References

[1] *Play around with ES6 features online http://www.es6fiddle.net/*

[2] *Tail call optimisation through arrow functions httpshttps://https://leanpub.com/understandinges6/read#leanpub-auto-arrow-functions*

[3] *Kangax support matrix http://kangax.github.io/compat-table/es6/*

By: Janardan Revuru

The author works at Hewlett Packard Enterprise R&D centre in Bengaluru. He is passionate about JavaScript and related Web frameworks. He can be reached for career consulting at *janardan.revuru@gmail.com*

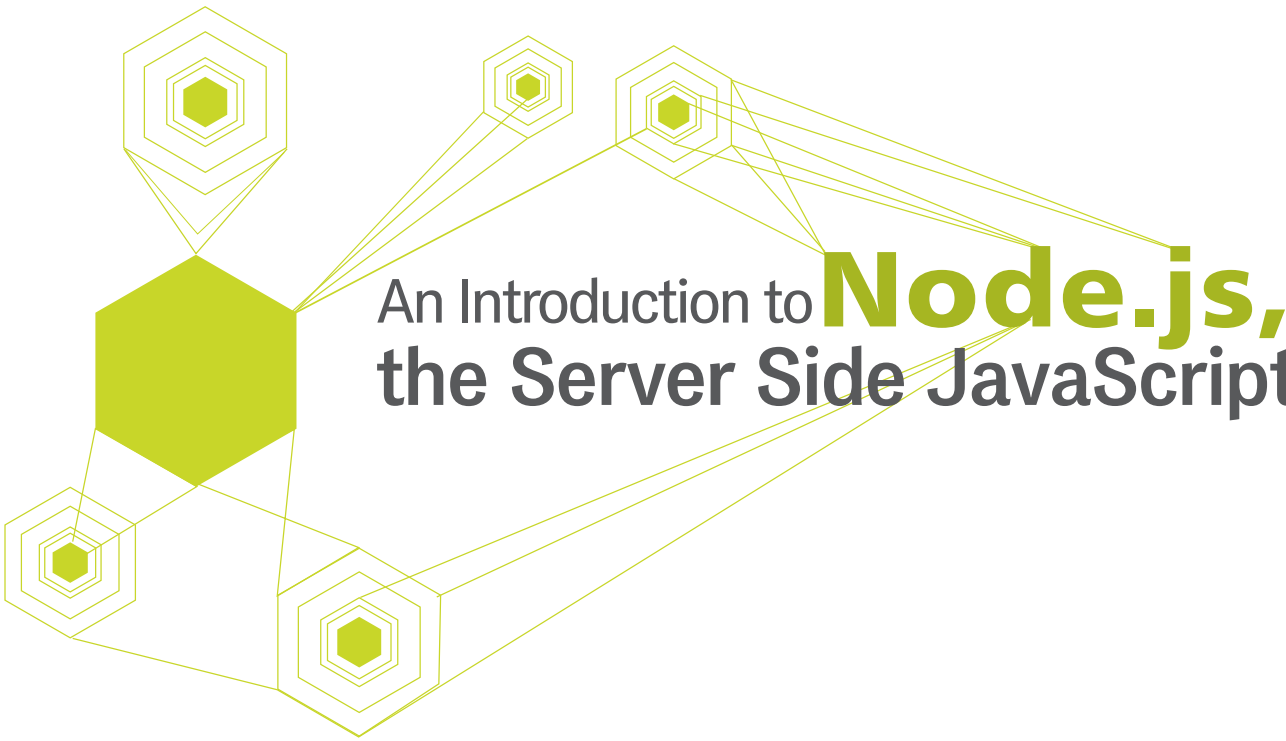

An Introduction to Node.js, the Server Side JavaScript

Node.js, often referred to as just Node, is a powerful tool that can run JavaScript applications on both the server side as well as the client side. Node.js can be used to write static file servers, Web application frameworks, messaging middleware, and servers for HTML5 multiplayer games. This article is a very elementary introduction to Node.js.

Node.js is a server side JavaScript built on Google's V8 JavaScript engine. It is an open source programming language that was developed by Ryan Dahl in 2009. It allows us to build scalable network applications, and is very fast when compared with other server side programming languages because it is written in C and the non-blocking I/O model. It is currently sponsored by Joynet, a software company specialising in high performance container native infrastructure. Node.js can run on Linux, Sun OS, Mac OS X and Windows platforms.

The steps below explain the difference between blocking code and non-blocking code models.

In the blocking code model, the steps are:

1. Read a file
2. Process the file
3. Print the result
4. Perform the next function

The steps in the non-blocking code model are:

1. Read a file
   1.1. When file reading is completed, process it
   1.1.1. When the processing is completed, print the result
2. Perform the next function

Internal architecture of Node.js

Node.js can be downloaded from *https://nodejs.org/* and is easy to install. Once installed in your system, you will see two of its main parts: the main node executable and the node package manager (NPM) executable. Node executable is simple, and you will pass the name of the main source file - for example, *node sample.js*.

The node executable will interpret the source code and execute it. Once it finishes, it will exit back to the shell. The NPM (node package manager) will add new packages to the node executable.

To install new modules, enter the command below:

```
npm install name-of-the-module
```

An HTTP server

If you are using PHP, a PHP file represents an HTML page. A Web server will accept a request and execute the code. But in Node.js, the main JavaScript file represents the entire Web server. It does not run inside a Web server.

Now let's save the code below as *main.js* and run it.

```
var http = require('http');
```

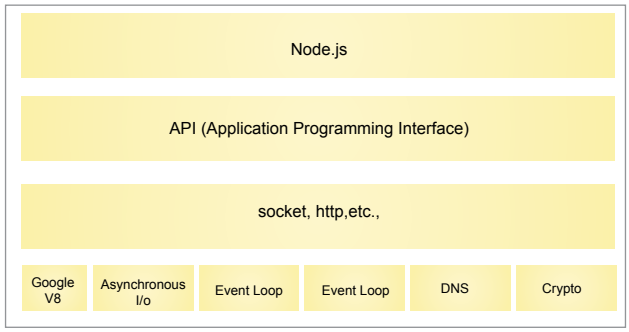

Figure 1: The internal structure of Node.js

```
var static = require('node-static');
var file = new static.Server();
http.createServer (function (request, response) {
            file.serve (request, response);
            response.writeHead(200, {'Content-Type': 'text/
plain'});
            response.end('Hello World!');
}).listen(1337, '127.0.0.1');
console.log('Server running at http://127.0.0.1:1337/');
```

How does the code given above work?

The global function *require()* makes a module available.

HTTP is a built-in module, whereas *node-static* is an external module. So you need to install the external module with the NPM, by using the following command:

```
npm install node-static
```

The keyword *new* is used to create an object.

After installing *node-static*, you will have a new folder called *node_modules* in your working directory. Now run the following command:

```
node main.js
```

After running the node, you need to open the Web browser and type *http://127.0.0.1:1337* into the address bar. You will see a simple Web page which says, 'Hello World!'

Creating your own module

It is very easy to create your own module in the Node.js environment.

Step 1: Save the code below as *custom_module1.js.*

```
exports.sayHello = function() {
            console.log('Hello, Welcome to Node.js!');
}
```

Step 2: Save the code below as *main.js.*

```
var hello = require('./custom_module1.js');
hello.sayHello();
```

Step 3: Run the command node *main.js.*

This is only an introductory article on Node.js. Interested readers are advised to explore further to find out about other features at *https://nodejs.org.* END

References

[1] *https://nodejs.org*

By: Nelson Joseph G

The author works at EinNel Technologies, Chennai, and can be contacted at nelson@einnel.com

OSFY Magazine Attractions During 2015-16

MONTH	THEME	BUYERS' GUIDE
March 2015	Open Source Firewall and Network security	SSD for Servers
April 2015	Web Development	Network Switches
May 2015	Virtualisation (containers)	Wireless Routers for SME
June 2015	Open source Databases	PaaS Solution
July 2015	Network Monitoring	MFD Printers for SMEs
August 2015	Mobile App Development	Hosting Solutions
September 2015	Backup and Data Storage	External HDD
October 2015	Programming Language	IaaS Solution
November 2015	Cloud Special	Firewall and UTMs
December 2015	Open Source on Windows	Online Backup solutions
January 2016	Android Special	Wifi Hotspot Devices
February 2016	Top 10 of Everything	External Storage

Create an *Exe* File for your Python Application in Windows

You can build your own search tool for Windows, using Python and allied technologies, by going through this tutorial. This tool will search for files which are located on a pen drive too. You can also hone your Python and programming skills by trying out this tutorial.

This article is for those with basic Python programming skills, looking to build simple Windows executable programs using Pywin32 and Pyinstaller. It doesn't provide solutions for building complex software programs, though.

There are a few prerequisites to creating your *exe* file, which are mentioned below:

- Windows OS
- Python 2.7.x
- Pywin32
- Pyinstaller

Python 2.7.x can be downloaded from *www.python.org*. You can get Pywin32 from *http://sourceforge.net/projects/pywin32/files/pywin32/*. Pyinstaller can be downloaded from *https://pypi.python.org/pypi/PyInstaller*.

In this article, we will try to cover the basics of creating a Windows executable with the help of Pywin32 and Pyinstaller. While installing Python, please select the option to add it into the class path variable of the system. After installing Python, open the command prompt window and type the following command:

```
pip install pyinstaller
```

Pywin32 is an executable file and can be installed easily.

After completing the installations, let's write a

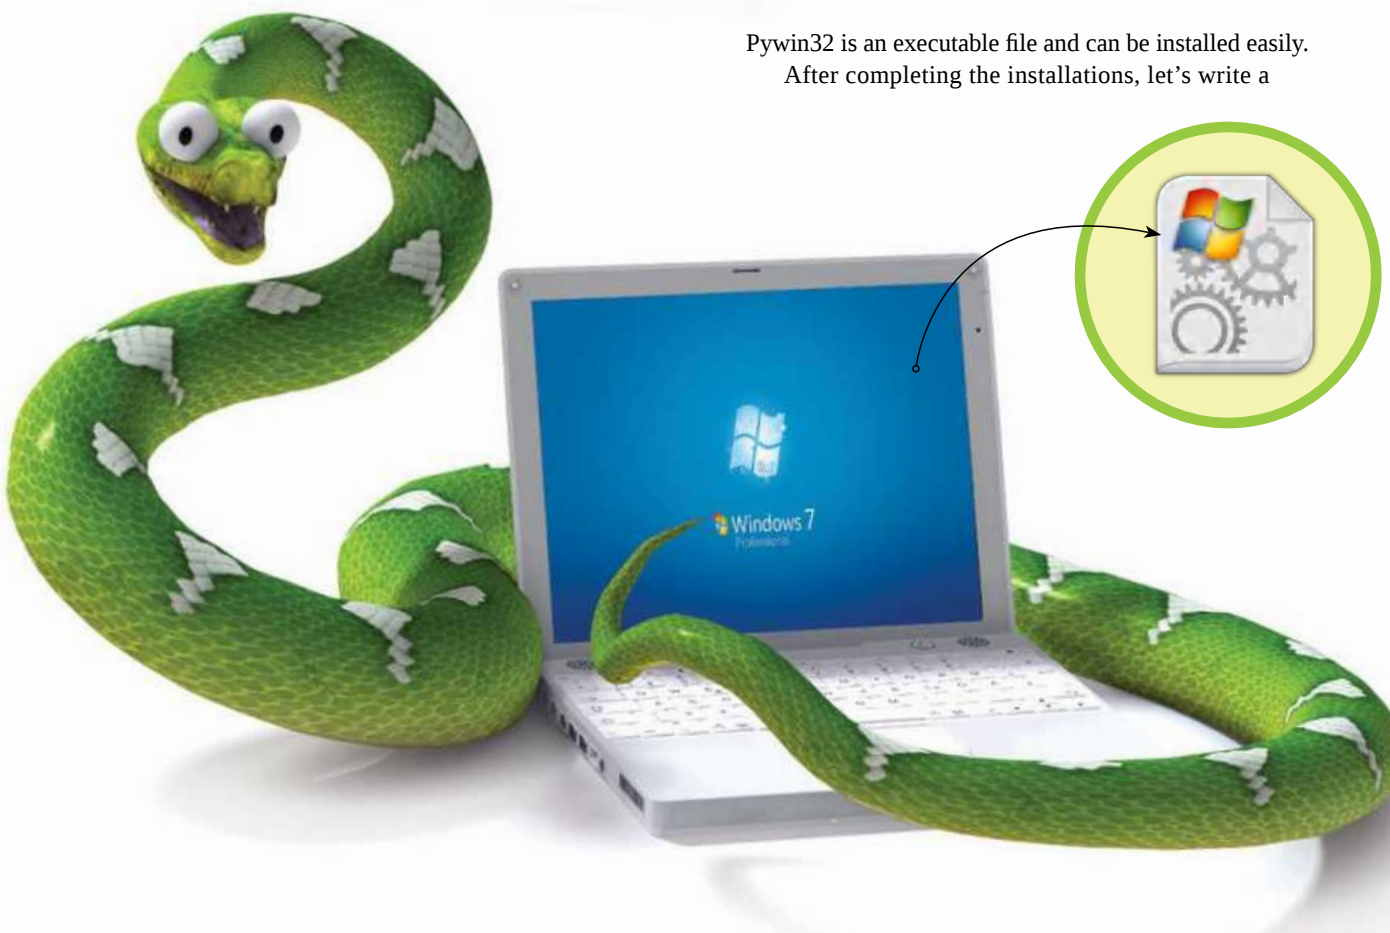

simple Python program file which will help to locate any file in the whole system, including external devices such as a pen drive. This Python file is built as an executable and placed in the Windows folder. The command *locate* is run with a combination of search strings such as the name of the file.

Given below is the code for the *Locate* program in Python:

```
Locate.py
import os
import re
import sys
from datetime import datetime

response = os.popen("wmic logicaldisk get caption")
list1 = []
for line in response.readlines():
              line = line.strip("\n")     # this line strip
the \n
              line = line.strip("\r")  # this line strip the
\r means carriage return
              line = line.strip(" ")  # use to
strip blank space
              if (line == "Caption" or line == ""):  # if
caption and "" appears this line ignore
                           continue
              list1.append(line)  # list1 now contain all
drives of windows
a = "\\"
def search1(list1):
              tf = 0
              td = 0
              for each in list1:      # take each drive
one by one
                           print "In ", each
                           for root, dir, files in
os.walk(each, topdown = True):   # this line return the root
dir (means start from drive),  sub driectory and files in the
root dir
                                        for file in files:
# start search from files
                                                  if
re.search(sys.argv[1].lower(), file.lower()):  # this line
match the substring given in command line argument with files
one by one

print root+a+file # if match found then make a path to print

tf = tf+1 # to count the files
                                        for d in dir:
                                                  if
re.search(sys.argv[1].lower(), d.lower()):
# this line for directories
```

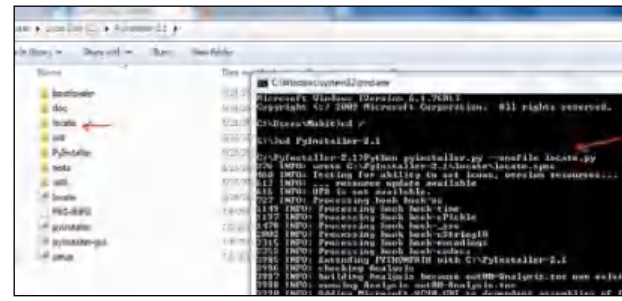

Figure 1: Output of *pyinstaller.py*

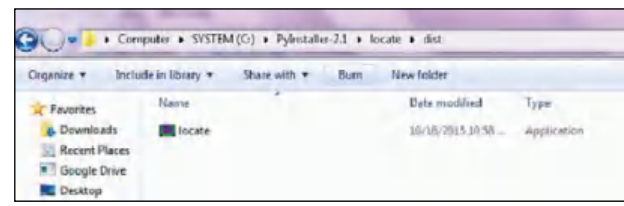

Figure 2: *Locate* command using the command prompt

```
print root+a+d  #print the path

td = td+1 # print the number of dirs matched
              print "Total File(s) are ", tf
              print "Total Directores are ", td
              print "\nThank you for using l4wisdom.com\n"

t1= datetime.now()
search1(list1)   # call the function.
t2= datetime.now()
total =t2-t1  # to calculate the total time taken in process
print "Searchingcompleted in " , total
```

Place the *locate.py* file in the Pyinstaller folder and open the command prompt. Go to the Pyinstaller folder and then execute the following command:

```
C:\PyInstall-2.1>Python pyinstaller.py --onefile locate.py
```

After running the command, you will get an output as shown in Figure 1.

Place the *locate.exe* file in the Windows folder of the system. This will make the system automatically pick the file through the Windows path variable. Now, go to the *run* command and provide the command *locate*, along with the file name you are looking for. If you don't remember the file name, then this command also takes the *search* query, as shown in Figure 2. END

By: Bhaskar Narayan Das

The author has over six years of experience in the IT industry and works as an application developer with IBM. He has worked with Java, Python and databases including NoSQL database related technologies. You can contact him at *bhaskarndas@gmail.com*

An App to Create Question Sets with Voice Input Answers

App Inventor empowers even newbies to create software applications for Android. Our regular readers have become quite familiar with App Inventor over the past few issues of *OSFY*. Noobs are advised to go through our back issues and catch up. This article is a tutorial on creating a quizzing app.

In this article let's create some popular content - a quiz with a set of questions and answers, along with an app that allows you to answer questions from a selected question set. We will also enable speech based inputs to answer questions, a novelty that could lead to other inventive applications too. Readers will require prior working knowledge of the *MIT AI 2* environment to benefit from this column.

The infrastructure

Let us examine what is needed as basic infrastructure to accomplish what we want to do. We are trying to create two apps: App1 to create data and App2 to use that data and enable some functionality. This means that data written by App1 needs to be communicated to or be available to App2. In the example featured in this article, we will use screens to simulate App1 and App2, and use TinyDB for communicating data between the apps. Refer to Figure 1. We could easily extend this using TinyWebDB and break the screens into totally independent apps. For this article, let's restrict ourselves to using TinyDB and a single app with multiple screens, which makes it easier to illustrate concepts.

Expanding on this concept further, the elements that we see in this article are explained in more detail in Table 1.

Element	Explanation
TinyDB	TinyDB is a non-visible component that stores data for an app. It is persistent and stores data in the form of tags and values. All apps running on the same system will share their TinyDB data store.
Screen	Screen is the basic container that shows a set of GUI elements in an app. A screen can invoke other screens and retain control.
SpeechRecognizer	SpeechRecognizer is a media component that allows the developer to interact with and use voice input functionality available in Android systems.
HorizontalArrange-ment	This is a GUI formatting element that allows items to be shown in one line or row, arranged from left to right.
VerticalAlignment	This is a GUI formatting element that allows items to be shown one below the other in the same column, arranged from top to bottom.

Table 1: Palette elements used

Next, let's look at how the design for the TinyDB tags is arrived at. TinyDB is a simple tag <-> value store, which has a simple interface to write a tag and its value into the database. It then allows you to fetch a value based on a tag name. Now, coming to the design of tags for our 'Question and Answer' system, the idea here is to identify a set of questions with a unique text based identifier called a QuestionSet. This comprises multiple questions and their answers. Hence, we need to know how many questions are present in a question set. We also need a way to identify all the QuestionSets already created in the system. So, we use the tags as explained in Table 2 for storing and retrieving the values needed on our app screens.

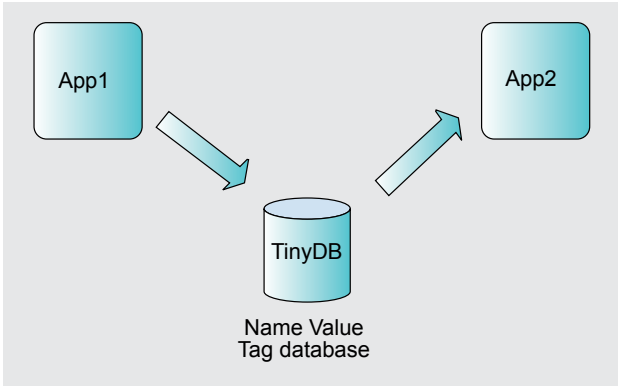

Figure 1: Apps communicate using TinyDB

Tag name	Explanation
QuestionSetList	A comma-separated list of all QuestionSets created using this app; e.g., set1, set2, set3, etc
<QuestionSet><n>Q	A text field representing the nth question in <QuestionSet>
<QuestionSet><n>A	A text field representing the answer to the nth question in <QuestionSet>
<QuestionSet>nTotalQ	A number field representing the total number of questions in <QuestionSet>

Table 2: The design of tags

The main screen

Now we come to the main screen. In each of these screens we first take a look at the elements we use from the palette in the 'design' view and what their names and values are to be. We will then present the code in the 'blocks' view and explain to an extent the role played by each code piece. In the main screen, which could not be renamed from 'Screen1' as far as I could make out, we see the elements described in Table 3 with their properties. Also note that TinyDB is a non-visible component. The screen itself should look like what is shown in Figure 2.

Name	Element type	Palette group	Property	Value
Screen1	Screen	NA	Title	Create or answer questions
CreateQ	Button	User interface	Text	Create questions
An-swerQ	Button	User interface	Text	Answer questions
ClearAll	Button	User interface	Text	ClearAll
Debug	Button	User interface	Text	Debug
TinyDB1	TinyDB	Storage		

Table 3: Screen1 elements

Ideally 'Create Questions' and 'Answer Questions' can be separate apps, as you may want to offer a different service to creators of questions and provide a different app for answering. But for this article, it would be fine to use the common app with different screens to be able to illustrate how they work using TinyDB for storage and retrieval of data. The blocks for Screen1 look like what's shown in Figure 3.

The main functionality of Screen1 is only to provide an interface to control invocation of the other screens, namely CreateQuestions, AnswerQuestions and also a DebugScreen. There is one more entry to ClearAll, which basically cleans up TinyDB of all the entries. Both the ClearAll and the DebugScreen aid in debugging during development. They are not expected to play this role in a production level app. The new procedure call, which opens another screen *screenname*, is available from the *Built-in→ Control* group of blocks. It basically takes a screen name as an argument and instantiates that particular screen. As of now, it looks like using the *Back* button is necessary to come back to the main screen and to do a clean-up in the child screen. We will come to that later in the next section.

Note: It is interesting to note that the screens cannot share any variables or any other common component. Also, even a copy-paste does not work across screens, either in design mode or in blocks mode.

The CreateQuestions screen

Now we come to the CreateQuestions screen. The elements of the screen are described in Table 4. The screen itself is shown in Figure 4.

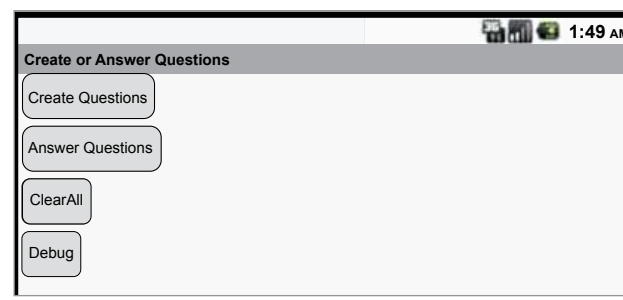

Figure 2: The main screen

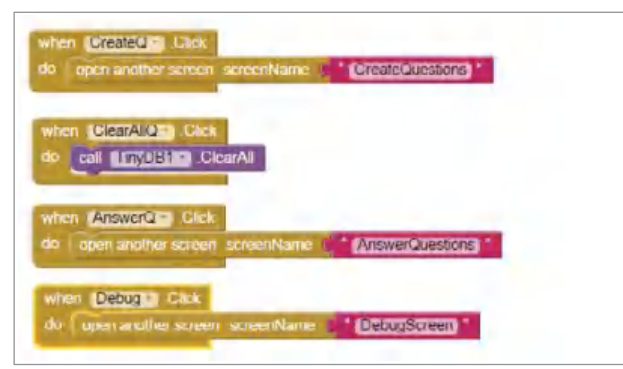

Figure 3: Blocks for main screen

Name	Element type	Palette group	Property	Value
CreateQuestions	Screen	NA	Title	CreateQuestions
HorizontalArrangement1	HorizontalArrangement	Layout	Width	FillParent
Label1	Label	User interface	Text	'Questions Set:'
QpaperName	TextBox	User interface	Hint width	'Name of Question Paper' FillParent
HorizontalArrangement2	HorizontalArrangement	Layout	Width	FillParent
Label2	Label	User interface	Text	'Questions'
Question	TextBox	User Interface	Hint Width Multiline	'Type the Question' FillParent checked
HorizontalArrangement3	HorizontalArrangement	Layout	Width	FillParent
Label3	Label	User interface	Text	'Answer'
QPaperName	TextBox	User interface	Hint width	'Type the Answer' FillParent
HorizontalArrangement4	HorizontalArrangement	Layout	Width	Automatic
SaveQnA	Button	User interface	Text	Save Question
GoBack	Button	User interface	Text	'Done (Go Back)'
TinyDB1	TinyDB	Storage		

Table 4: CreateQuestions' screen elements

TinyDB has some very simple interfaces for CRUD (create, read, update and delete) operations. In fact, create and update are one and the same operation as the presence of a tag before them does not make any difference. In addition, a GetTags interface allows one to read all the tags present in the database. The StoreValue interface allows one to store a value against a particular tag. The GetValue interface allows the reading of the value stored against a particular tag. There is also a ClearTag interface, which deletes a tag and its value from the database. Putting these together we have the blocks for the CreateQuestions screen, as shown in Figure 5.

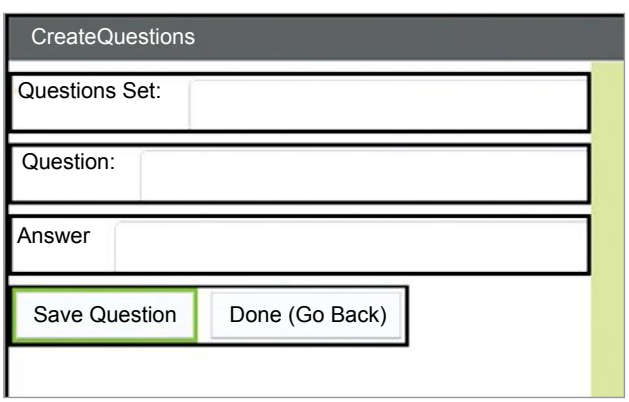

Figure 4: CreateQuestions UI

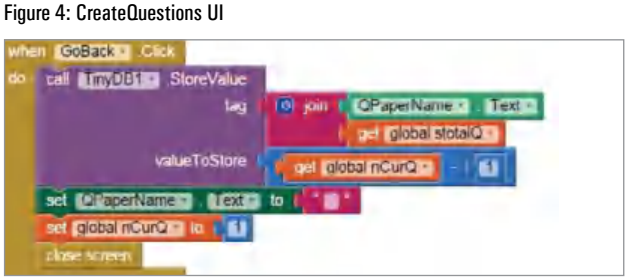

Figure 5: CreateQuestions blocks

AnswerQuestions screen

Now we come to the AnswerQuestions screen. The elements here are similar to CreateQuestions, except that we add an element for voice input, called SpeechRecognizer. Also, we create an image button to indicate voice input is enabled. To do this, we can upload the image of a microphone into the project's media section. In the current case, I have uploaded a file *microphone.png* into the project. Then we can set the image for the button. Next, let's add a couple of buttons to navigate forward and back, amongst the questions. We describe elements of the screen in Table 5. The screen itself is shown in Figure 6.

Note: There are three values for width, height, etc. These could be automatic, which means that they will get adjusted based on their properties, like a text box getting adjusted based on text input. It could also be 'Fill Parent', which means that it will fill as much as the parent element allows it to; e.g., for our 'Mike' button, we want its height to fill as much as the parent will allow. Automatic will make it too big, based on the image size. We can also make it a percentage of the available space and, in this case, you can adjust the percentage out of 100 per cent to different elements available.

The blocks for AnswerQuestions in Figure 7 are similar to those of CreateQuestions except for additional checks and controls over the following items.

Validating the answer: Once the answer is entered, we need to validate it against the answer stored for the question when it was created. While doing this, I decided to show an incorrect answer by changing the colour of the answer to red. When the correct answer is entered, it changes back to black.

SpeechRecognizer input: When the 'Mike' button

Name	Element type	Palette group	Property	Value
AnswerQuestions	Screen	NA	Title	AnswerQuestions
HorizontalArrangement1	HorizontalArrangement	Layout	Width	FillParent
Label1	Label	User interface	Text	'Questions Set:'
QpaperName	TextBox	User interface	Hint width	'Name of the Question Paper' FillParent
VerticalArrangement2	VerticalArrangement	Layout	Width	FillParent
Label2	Label	User interface	Text	'Question'
Answer	TextBox	User interface	Hint Height Width Multiline	'Type the Answer' Automatic 85% checked
Mike	Button	User interface	Image Height Width Multiline	Microphone.png FillParent 15% " "[empty string]
HorizontalArrangement4	HorizontalArrangement	Layout	Width	Automatic
Previous	Button	User interface	Text	Previous
Next	Button	User interface	Text	Next
GoBack	Button	User interface	Text	Done (Go Back)
TinyDB1	TinyDB	Storage		

Table 5: AnswerQuestions screen elements

element is clicked, we invoke the speech input component. We set the value returned from this to the 'Answer' text box.

Controlling the current question with the previous and the next one: We use a question tracker as a number starting at 1, in order to navigate between questions. This needs to be changed and the corresponding question loaded into the screen elements. Also, to do this neatly, we create a procedure loadQuestion, which reads the current question number and loads the question accordingly.

Debugging

During development, I wanted to understand how the tags were created and retrieved. Thus I felt the need to have a debugging interface with which I could walk through the tags and values already created. The DebugScreen is meant for this and is created with simple elements as shown in Table 6.

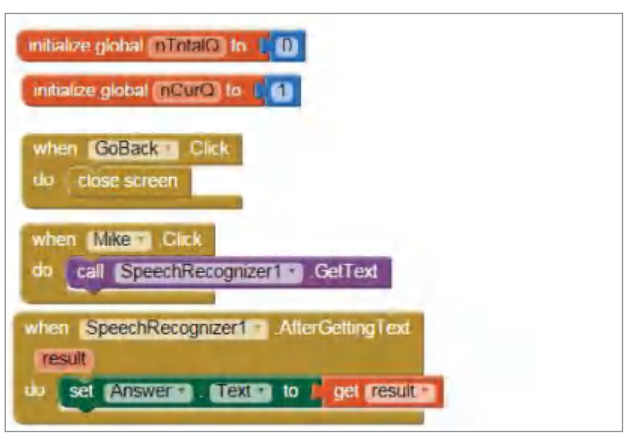

Figure 6: AnswerQuestions screen

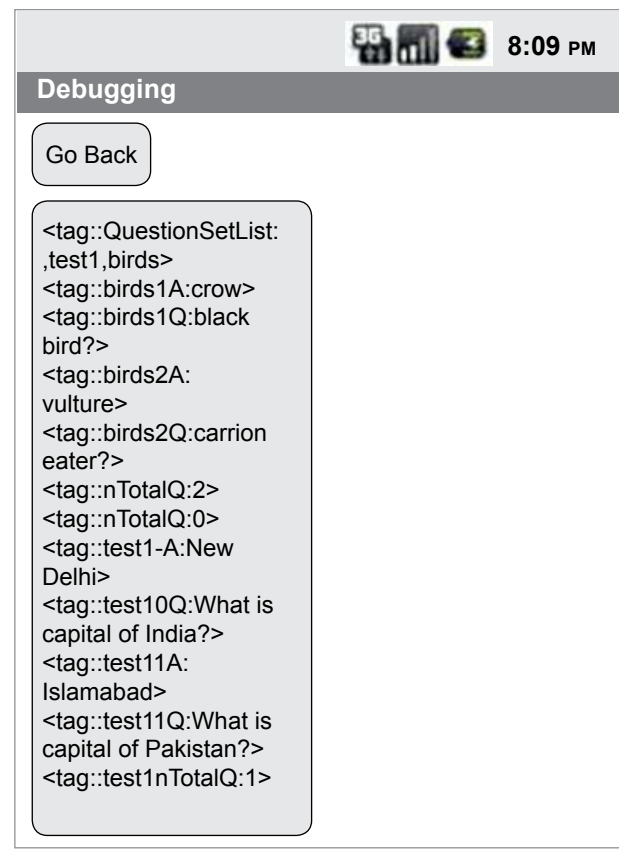

Figure 7: AnswerQuestions screen blocks

The DebugScreen is shown in Figure 7.

The blocks code for DebugScreen looks like what is shown in Figure 8. For operations that you want to perform

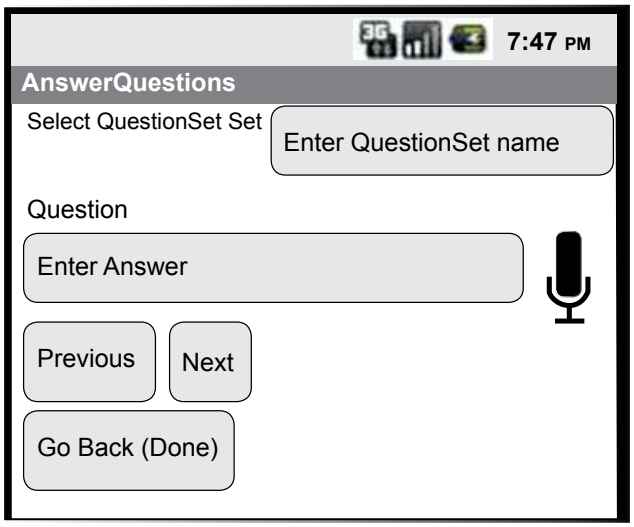

Figure 8: DebugScreen output

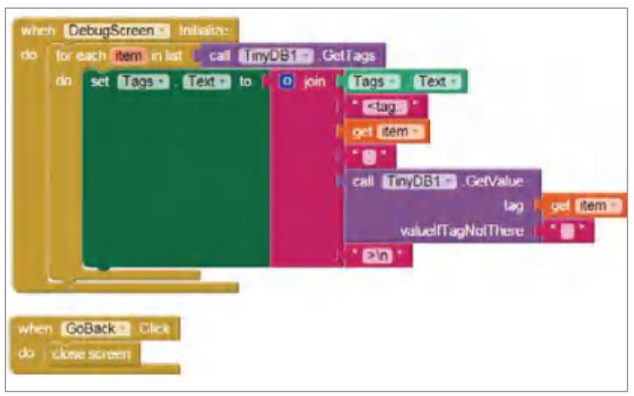

Figure 9: DebugScreen blocks

during screen load, you need to implement the Screen. Initialise method and implement your functionality.

Important points for debugging and more

The connect to the 'AI Companion' option does not work when a VPN network is connected. The error is not usable, and it occurred to me later that VPN is interfering with the connectivity between the AI2 in the cloud and the mobile phone.

The speech recogniser does not work in the Emulator mode and gives an error when invoked. It looks like there is some error in the settings, even though my laptop is capable of microphone input.

Every environment in which you run sets up its own TinyDB database and, hence, when you test freshly in Emulator or in AI Companion you will find that the QuestionsList needs to be set afresh.

TinyDB on your phone or emulator is common across all apps and, hence, may have tags and values that are not intended for your app and may even interfere with your data.

Name	Element type	Palette group	Property	Value
Debug-Screen	Screen	NA	Title	Debugging
GoBack	Button	User interface	Text	GoBack
Tags	TextBox	User interface	Text	'Questions Set :'

Table 6: DebugScreen elements

What next?

This app is only a starter. You could consider doing the following to improve it:

1. Use TinyWebDB to store and retrieve data across apps.
2. Use microphone input to also create questions.
3. Make the app completely media-friendly by including a SpeakQuestion option too, to read out loud.
4. Include the ability to give multiple options as answers. *Hint:* You can store an xml in the answer for a tag. END

By: T.V. Krishnamurthy

The author has more than 15 years of experience in the telecom industry. He has been actively involved in radio technologies for nearly a decade while working with Siemens Communications, NSN, Radisys and now Nokia Networks. He lives and works out of Bengaluru in India.

DOT

A Language that Helps You to Draw Graphs

DOT is a simple, plain text language that can describe graphs which both humans and computers can use. It can be used to generate graphs in different formats such as *.ps* and *.pdf*. This article describes how to install DOT, draw a graph, and then generate a graph that depicts an execution flow of a program in C.

DOT is a language you can use to textually represent a graph, so that it can be processed by the dot tool to render the graph as a graphical representation in different formats like *.ps*, *.pdf*, etc. It has been developed at AT&T Labs as part of the Graphviz project, which is a collection of tools for graph visualisation. DOT is open source free software published under the Eclipse Public License.

Installation

For the purposes of this article I will use the Ubuntu 14.04 LTS OS version to install DOT; however, you can use any distribution you choose. The dot tool comes with the Graphviz package along with many other graph drawing tools. Run the following command to install Graphviz:

```
$ sudo apt-get install graphviz
```

Understanding graphs

Before we start drawing our first graph, let's first understand what it is, in a non-technical way. A graph G consists of a set of nodes called vertices and a set of edges connecting a pair of vertices. If the edges have direction, it is a directed graph; else, it is an undirected graph. So the diagram in Figure 1 is a directed graph or digraph. v1, v2, v3 and v4 are vertices, and e1, e2 and e3 are edges connecting the ordered vertex-pairs (v1, v2), (v1, v4) and (v4, v3), respectively. All edges have a direction.

Our first graph using DOT

Open a text editor and enter the following code to draw the graph as shown in Figure 1; save it as *firstgraph.dot*.

```
digraph G {
      v1 -> v2 [label="e1"];
      v1 -> v4 [label="e2"];
      v4 -> v3 [label="e3"];
}
```

To generate the graph in PDF format, run the following:

```
$ dot firstgraph.dot -Tpdf -o firstgraph.pdf
```

Run the following command to see the graph (Figure 2).

```
$ evince firstgraph.pdf
```

Generate an execution flow graph

Now that we know how to draw basic graphs, let's generate a graph for the execution flow of a C program. The following C program *oddeven.c* checks whether a given integer number is odd or even.

```
#include<stdio.h>
int main()
{
```

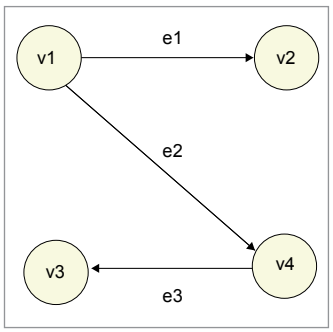

Figure 1: A simple directed graph

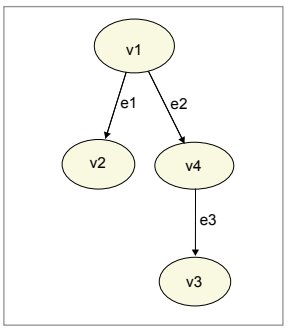

Figure 2: A DOT generated graph

```
    int num;

    printf("\nEnter an integer:");
    scanf("%d",&num);

    if( num%2==0 )
    {
      printf("%d is even number\n",num);
    }
    else
    {
      printf("%d is odd number\n",num);
    }

    return 0;
}
```

This program is so simple, it needs no explanation. We will instrument this program to automatically generate a *.dot* program file for the execution flow, based on a particular input. Shown below is the instrumented code *oddevendot.c*

```
#include<stdio.h>

int main()
{
    int num;
    FILE *fp;

    printf("\nEnter an integer:");
    scanf("%d",&num);

    if( (fp=fopen("oddeven.
dot","w"))!=NULL )
    {
      fprintf(fp,"digraph G
{\n");

      fprintf(fp,"\tstart
[label=\"%d%%2 == 0 ?\"]\n",num);
      if( num%2==0 )
      {
        fprintf(fp,"\tstart -> end [label=\"true\"]\n");
        printf("%d is even number\n",num);
        fprintf(fp,"\tend [label=\"EVEN\"]\n");
      }
      else
      {
        fprintf(fp,"\tstart -> end
[label=\"false\"]\n");
        printf("%d is odd number\n",num);
        fprintf(fp,"\tend [label=\"ODD\"]\n");
      }
      fprintf(fp,"}\n");
      fclose(fp);
    }
    else
    {
      printf("\nfopen() failed\n");
    }

    return 0;
}
```

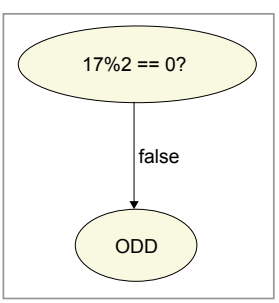

Figure 3: Execution flow graph for input 17

This program generates the dot file *oddeven.dot*. We will create the PDF graph from this file using DOT. Let's try to run, compile, execute and create the graph as shown below:

```
$ gcc oddevendot.c -o oddevendot
$ ./oddevendot
Enter an integer:17
17 is odd number
$ dot oddeven.dot -Tpdf -o oddeven.pdf
$ evince oddeven.pdf
```

Figure 3 shows the graph that will be displayed. Though this is overkill for such a simple program, it is only to demonstrate the usefulness of the programmatic way of drawing graphs.

I have only covered a very few features of the DOT language. Please explore the resources to find out how powerful and flexible the language is. You can also explore the 'neato' tool Graphviz package for drawing undirected graphs. END

References

[1] *http://www.graphviz.org/*
[2] *https://reference.wolfram.com/language/ref/format/DOT.html*

By: Ramaprasad Chakraborty

The author works at ISRO, Sriharikota, as scientist engineer SD and has six years of industry experience. He is interested in many aspects of computer engineering and loves to explore different open source technologies.

Playing Games on GNU Emacs

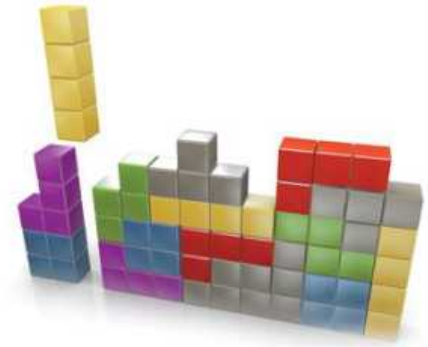

In this article in the GNU Emacs series, let's learn how to play games and try out some software for fun.

Ten games that are available and can be played easily in GNU Emacs are listed and discussed in this article. We start with the most popular one, Tetris.

1. Tetris

Tetris is available in GNU Emacs and you can launch it using *M-x tetris*. The 'Up' arrow key can be used to spin the tile, while the space key can quickly drop the tile. You can pause the game by pressing the 'p' key. At the end of the game, a score is displayed in a buffer. For example:

```
112 shakthi shakthi <author@shakthimaan.com>  Mon Nov  2
21:12:04 2015
52 shakthi Shakthi Kannan <author@shakthimaan.com>  Mon Nov
2 15:24:43 2015
11 shakthi shakthi <author@shakthimaan.com>  Mon Nov  2

15:29:39 2015
```

A screenshot of Tetris in GNU Emacs is shown in Figure 1.

2. Pong

You can play Ping Pong in GNU Emacs with two players, using *M-x pong*. One player can use the 'Up' and 'Down' arrow keys, and the second player can use the number keys – '4' and '6', respectively. Each paddle is represented by three square tiles. If the ball hits the edge of the paddle, it changes its direction. You can use the 'p' key to pause the game. A live score is displayed during the game. A screenshot of Pong is shown in Figure 2.

3. Doctor

GNU Emacs comes with a psychotherapist with whom you can chat. Simply type *M-x doctor* and it will start a session with the therapist. After you type in your statement, you need to hit *Enter* twice, and the therapist will reply. It will usually ask you another question, and you can continue to chat with the doctor. An example of a session is given below:

```
I am the psychotherapist.  Please, describe your problems.
Each time you are finished talking, type RET twice.
Hi How do you do?  What brings you to see me?
I am trying to talk to the doctor. Why do you say you are
trying to talk to the doctor?
I want to share this transcript with my readers.
What makes you believe you might want to share this
transcript with your readers?
I can, from GNU Emacs, copy this text and paste it in my
article
Emacs?  Hah!  I would appreciate it if you would continue.
I am looking forward to watching today's cricket match
that India is playing.
Is it because India is playing that you came to me?
Yes, who do you think will win?
Earlier you said you would like to share this transcript
with your readers?
Yes, I am still going to do that.
Can you elaborate on that?
I am writing a GNU Emacs series for the OSFY magazine.
Are you sure?
Yes, I have been writing for the past few months.
Maybe your life has something to do with this.
I want people to learn and use GNU Emacs to improve their
productivity.
Are you sure about that?
Absolutely!
Perhaps you could continue.
```

4. Bubbles

The objective of the Bubbles game is to remove a series of connected bubbles in the fewest number of moves. The game can be started using *M-x bubbles*. As you move through the grid of bubbles, the connected bubbles with the same colour get highlighted. You can press the space bar to remove them. A screenshot of Bubbles is provided in Figure 3.

5. Conway's Life

The shortcut *M-x life* will start the Conway's Game of Life in a new GNU Emacs buffer. The population increases in the course of time, and the generation count is shown in the mode line. A screenshot of the cellular automation running inside GNU Emacs is provided in Figure 4.

6. Morse and NATO conversions

You can convert text to Morse code and vice versa by selecting the text in a buffer and using *M-x morse-region*. For example, the text 'morse code' gets converted to the following:

```
--/---/.-./.../. -.-./---/-../.
```

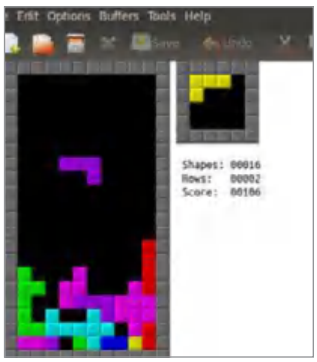

Figure 1: Tetris

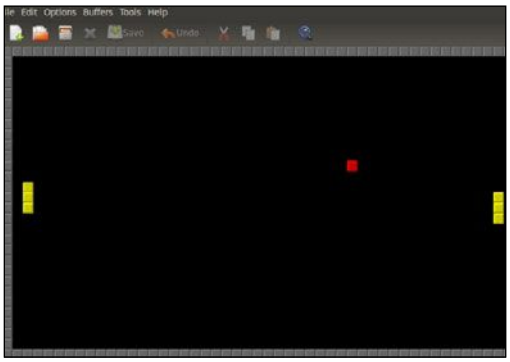

Figure 2: Pong

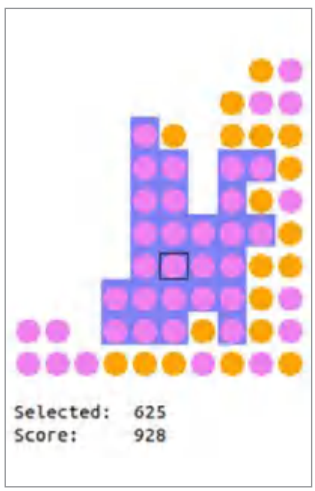

Figure 3: Bubbles

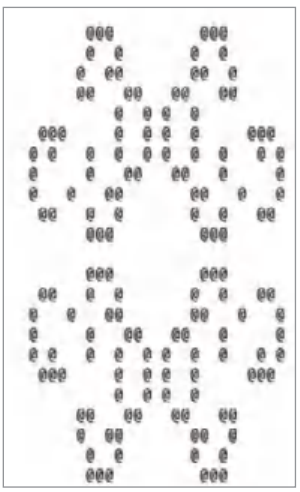

Figure 4: Conway's Life

You can get back the text by selecting the Morse text and applying *M-x unmorse-region*. Similarly, if you have a word that you would like to spell using the NATO phonetic alphabet, you can use *M-x nato-region*. To convert it back, you need to use *M-x denato-region*. For example, the text 'abc' gets converted to:

```
Alfa-Bravo-Charlie
```

7. Snake

The Snake game can be started using *M-x snake*. You can use the arrow keys to move the head. As you play, red boxes appear in the window. If you go over them, the length of the snake increases along with the score. At the end of the game, a summary of the scores is shown. A screenshot of the Snake game is shown in Figure 5.

8. Dunnet

Dunnet is a text based adventure game that needs to be started in batch mode as shown below:

```
$ emacs -batch -l dunnet
Dead end
You are at a dead end of a dirt road.  The road goes to the
east.
In the distance you can see that it will
eventually fork off.  The
trees here are very tall royal palms, and
they are spaced equidistant
from each other.

There is a shovel here.
```

The *help* command gives you the context of the game:

```
>help
Welcome to dunnet (2.01), by Ron Schnell (ronnie@driver-aces.
com).
Here is some useful information (read carefully because there
are one
or more clues in here):
- If you have a key that can open a door, you do not need to
explicitly
  open it.  You may just use 'in' or walk in the direction of
the door.
- If you have a lamp, it is always lit.
- You will not get any points until you manage to get
treasures to a certain
  place.  Simply finding the treasures is not good enough.
There is more
  than one way to get a treasure to the special place.  It is
also
  important that the objects get to the special place
*unharmed* and
  *untarnished*.  You can tell if you have successfully
transported the
  object by looking at your score, as it changes immediately.
Note that
an object can become harmed even after you have received
points for it.
If this happens, your score will decrease, and in many cases
you can never
get credit for it again.
- You can save your game with the 'save' command, and restore
it
  with the 'restore' command.

- There are no limits on lengths of object names.

- Directions are: north, south, east, west, northeast,
southeast, northwest,
                 southwest, up, down, in, out.

- These can be abbreviated: n, s, e, w, ne, se, nw, sw, u, d,
in, out.

- If you go down a hole in the floor without an aid such as a
```

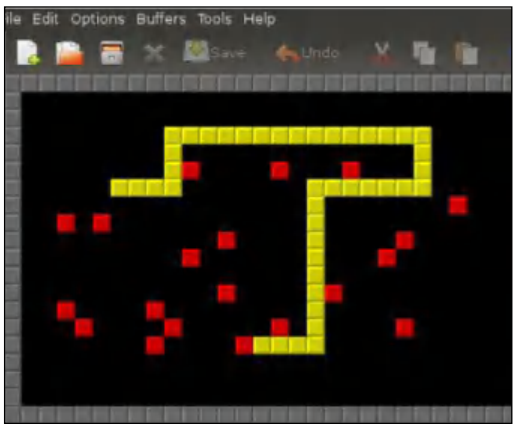

Figure 5: Snake

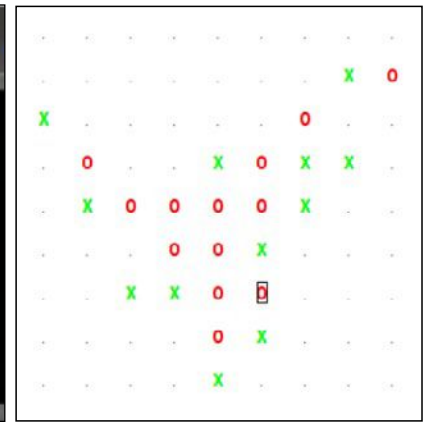

Figure 6: Gomoku

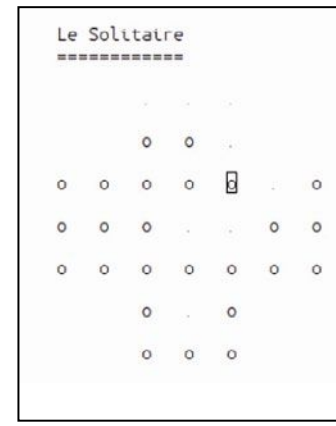

Figure 7: Le Solitaire

```
ladder,
  you probably won't be able to get back up the way you came,
if at all.

- To run this game in batch mode (no Emacs window), use:
    emacs -batch -l dunnet
NOTE: This game *should* be run in batch mode!

If you have questions or comments, please contact ronnie@
driver-aces.com
My home page is http://www.driver-aces.com/ronnie.html
```

You can then give directions and proceed with the game. An example of a session is shown below:

```
>n
You can't go that way.
>s
You can't go that way.
>e
E/W Dirt road
You are on the continuation of a dirt road.  There are more
trees on
both sides of you.  The road continues to the east and west.
There is a large boulder here.
>in
You can't go that way.
>w
Dead end
There is a shovel here.
>d
You can't go that way.
>e
E/W Dirt road
There is a large boulder here.
>n s e
You can't go that way.
>in
You can't go that way.
>e
Fork
You are at a fork of two passages, one to the northeast, and
one to the
southeast.  The ground here seems very soft. You can also go
back west.
>w
E/W Dirt road
There is a large boulder here.
```

You can exit the game by typing 'quit' at the prompt (>).

9. Gomoku

Gomoku is a strategy board game for which you need to get five pieces in a row (any direction) to win. You and the computer will take turns to play the game. You need to use the *Enter* key to mark your cross in a position, and the computer will mark its position with a circle. The game ends when either player gets five continuous pieces. A screenshot is shown in Figure 6.

10. Le Solitaire

Le Solitaire is a strategy game that consists of stones represented by 'o' and holes represented by '.' (a dot). The objective of the game is to remove all the stones except the last one. You can jump over another stone to create a hole. You can use the *Shift* key with the arrow keys to move a stone. A screenshot of the game in progress is shown in Figure 7.

All the games and examples were tried on GNU Emacs 24.5.1. END

By: Shakthi Kannan

The author is a free software enthusiast and blogs at *shakthimaan.com*.

Getting Started with Theme Development in WordPress

A WordPress theme determines the look and basic functions of a WordPress site. There are hundreds of themes available for download—free or on payment. The WordPress website itself has a number of free themes. But when one needs a special theme for a particular need, then it's time to develop one's own theme. This article will guide the reader through the theme development process.

WordPress is a PHP and MySQL based free and open source CMS (content management system). Used in more than 23 per cent websites, it's the most powerful blogging and website content management system in existence today. Famous blogs like Mashable and TechCrunch are both on WordPress. News outlets like 'The New York Times' blogs and CNN's on-air personality blogs all use WordPress, too. It is free and can be downloaded for self-hosted installations from WordPress.org.

Themes in WordPress

A WordPress theme provides all of the front-end styling for your site. These themes provide:

- The overall design or style of the site
- Font styling
- Colours
- Widget locations
- Page layouts
- Styles for blog posts and blog archives

Naming a WordPress theme

To build a theme in WordPress you have to be aware of its directory structure. First, create a sub-folder in the *wp-content/themes* directory in your WordPress folder, e.g.,

my_theme. The name of the folder should correspond with the name of the theme you want to create and should be unique in the theme folder.

To develop a theme in WordPress, you must have the HTML document ready with you. Place this HTML file with its assets (CSS and JS) in the theme directory.

Breaking up the HTML template into PHP files

The next step involves creating four PHP files. Using the HTML file, we have to create *index.php*, *header.php*, *sidebar.php* and *footer.php* by breaking the HTML file into four parts and saving them in the same theme folder. Just create one more file in the same directory *style.php* for styling purposes.

How these files work is shown below.

header.php: This file will contain the code for the header section of the theme.

index.php: This is the main file for the theme. It will contain the code for the 'Main Area' and will specify where the other files will be included.

sidebar.php: This file will contain the information about the sidebar.

footer.php: This file will handle your footer.

style.css: This file will handle the styling of your new theme.

Now let's take a closer look at each file and what it should contain.

header.php

```
<html>
<head>
<title>My Theme</title>
<link href="<?php echo get_template_directory_uri(); ?>/
css/header.css" rel="stylesheet" type="text/css" />
<script language="javascript" type="text/javascript"
src="<?php echo get_template_directory_uri();  ?>/jquery-
1.8.3.js"></script>
</head>
<body>
<div id="wrapper">
<div id="header">
<h1>HEADER</h1>
</div>
```

index.php

```
<?php get_header(); ?>
<div id="main">
<div id="content">
<h1>Main Area</h1>
</div>
<?php get_sidebar(); ?>
</div>

<?php get_footer(); ?>
```

sidebar.php

```
<div id="sidebar">
<h2 >Sidebar Area</h2>
</div>
```

footer.php

```
<div id="footer">
<h1>FOOTER</h1>
</div>
</div>
</body>
</html>
```

We can fetch other parts of our HTML template into *index.php* by calling a few functions like:

1. *get_header()* to call *header.php*
2. *get_footer()* to call *footer.php*
3. *get_sidebar()* to call *sidebar.php*
4. *get_template_directory_uri():* This function is used to provide the full path of the theme directory to access the style sheet and scripts.
5. *the_title():* This function is used to show the title of the post.

To fetch the post in our HTML theme, we have to use the WordPress predefined function:

```
<?php if (have_posts()) : while (have_posts()) : the_
post();  ?>
<h1><?php the_title(); ?></h1>
<h4>Posted on <?php the_time('F jS, Y') ?></h4>
<p><?php the_content(__('(more...)')); ?></p>
<hr> <?php endwhile; else: ?>
<p><?php _e('Sorry, no posts matched your criteria.'); ?></
p><?php endif; ?>
```

It will fetch all the contents and the titles of the posts, along with the time when these were posted. We have to place the above code in our HTML template, accordingly. END

By: Gaurav Kakran

The author is senior associate, Web development at Fierce Hound Media Pvt Ltd, Gurgaon. He is fond of developing Web applications in PHP, and can be reached at *gauravkakran7@gmail.com*

How to Run Linux on a Windows Machine

Kipling once said, "Oh, East is East and West is West, and never the twain shall meet ..." This same sentiment was often echoed about Linux and Windows. However, today, Linux can run on a Windows machine. This article discusses the several ways in which this is possible.

When you meet any Windows user and talk about using Linux, the first thing you will probably hear is, "I cannot use Linux because I don't know the commands and I am not a computer guru." The fact, however, is that today's top Linux desktop distributions such as Mint, openSUSE and Ubuntu are much easier to use than Windows 8.

Can Linux be run on Windows?

Another misconception among people who want to use Linux but are still on Windows is that the only way to use Linux and Windows on the same hardware is by making it dual boot. And they feel that with Windows 8 bringing in secure boot, a dual boot could turn into a nightmare.

But the truth is that it is not only possible but also very easy to start using Linux on a running Windows system, and there are many options for doing this.

Ways to use Linux on Windows

Let's look at the various options available to run Linux on Windows that we will be discussing in this article.

1. Take the online Ubuntu tour
2. Try Linux CLI on the Web browser
3. Try Linux-like Windows software
4. Try Linux with a live CD
5. Try Linux using a pen drive
6. Try Linux using a virtual machine
7. Try coLinux based distributions

Now let's start discuss each option in detail.

Take the online Ubuntu tour

This should be one of the easiest ways to try your hand at Linux. Ubuntu offers a Web application that simulates an installed Ubuntu desktop in your Web browser. Just open any Web browser and go to *http://tour.ubuntu.com/en/* and start exploring the Ubuntu desktop version. This gives an interactive view of the Ubuntu desktop, where the user can experience various Ubuntu menu options.

Try Linux CLI on a Web browser

The earlier option allows you to experience the Ubuntu GUI. But if you want to try Linux in its original form, i.e., CLI, open any browser from Windows, go to *http://bellard.org/jslinux/index.html* and start trying out Linux CLI. This is a JavaScript based emulator which lets you run various Linux commands. You can write shell scripts, use regular expression utilities like SED and AWK, experience editors like Vi and Emacs, and can also look up the man pages.

Try Linux-like Windows software

Wouldn't it be nice if we could install Linux on a Windows machine just as we install any other software in Windows? Well, let's take a look at WUBI.

WUBI is the acronym for 'Windows based Ubuntu installer'. It is a very small utility that helps you install Ubuntu on a system running Windows. The important thing to note here is WUBI does not require a separate partition, whereas it creates a special file on the Windows partition and uses that file as the Ubuntu drive. The user can very well uninstall Ubuntu from the Windows control panel.

To install Ubuntu using WUBI, one needs to follow a few basic steps:

1. Download the WUBI utility from *http://www.ubuntu.com/download/desktop/windows-installer* and then install it on a Windows desktop machine.
2. Start the WUBI utility and in the *Installation directory*, select the Ubuntu flavour to install it from the drop down menu.
3. Click on *Install* to start the installation.
4. The normal Ubuntu installation will proceed and in the end, you will see the wizard asking you to reboot.
5. After reboot, the next boot user will see the Windows boot manager displaying two operating systems to boot from.
6. Choose Ubuntu, boot into it and start exploring it.

Try Linux with a live CD

Till now, we have looked at how to experience emulated versions of Linux on Windows. Let's move a step ahead and try out real Linux. Now, we certainly have to leave our Windows environment. But here, the objective is to experience Linux without affecting the already installed Windows in any way.

Let's see how we run Linux with a live CD/DVD.

1. First, download the desired Linux flavour LiveCD image from any of the hundreds of available repositories, some of which are listed below:
   a. *https://livecdlist.com/*
   b. *http://www.knopper.net/knoppix/index-en.html*
   c. *http://www.linux.com/directory/Distributions/livecd*
2. Now go to your systems BIOS setting and alter the boot options to enable you to boot from the CD/DVD.
3. We are now ready to boot into Linux. Just put the CD/DVD in your disk drive and reboot your desktop.
4. When the boot is complete, you will see that you have booted into Linux.
5. The only disadvantage of using a live CD is that nothing is persistent and everything starts from scratch with the next boot.

Try Linux using a pen drive

We have booted Linux with a CD/DVD image but what if our system doesn't have a disk drive? The solution is to boot into the Linux operating system from a USB stick. Any normal USB stick can be used to install a live Linux image, which gives us a mobile operating system ready to use anywhere, right in our pocket. To install Linux on a pen drive, you need one with a minimum capacity of 2GB, a Linux installer ISO image, and a computer to boot from.

There are broadly two ways of installing Linux on USB flash drives.

1. **Installing Linux on a USB flash drive manually:** Here, we start a normal Linux installation and at the stage where we have to select the disk drive on which to install, the user carefully selects the USB flash disk drive. The only concern with this approach is that the user must have a very basic understanding of how disk drives are represented in Linux, and should be able to differentiate between USB flash storage and other disk drives present in the system.
2. **Use the USB Linux installer utility:** The concern mentioned in relation to the first approach is not so serious. Various Linux distros provide a solution to it. There are hundreds of free utilities available on the Internet, which install Linux on a USB flash drive from Windows. A few such utilities are listed here.

a. ***LinuxLive USB Creator:*** This is a free, open source Windows application that allows you to install Linux on a flash drive. You can download and install the application first; and then, in just four to five steps, have your USB flash drive ready to boot Linux:
   i. Start the application from the Windows *Start* menu and choose the disk from the *drives* list.
   ii. Select the ISO file or a CD as the installation source.
   iii. Choose the size of persistent data, so that the contents on this disk space on the flash drive will persist on it even after reboot.
   iv. Now, click on the 'lightning' button to start the image creation/installation.
   v. This utility can be downloaded from *http://www.linuxliveusb.com/*.

b. ***Universal USB Installer:*** Universal USB Installer or UUI is another live Linux USB creator that allows you to install Linux on a flash drive from Windows. Here again, you just need to download and install the utility and follow a couple of very basic self-explanatory steps to have a Linux bootable USB flash drive ready. UUI can be downloaded from *http://www.pendrivelinux.com/universal-usb-installer-easy-as-1-2-3/*.

c. ***UNetbootin:*** This is another application that allows you to create a bootable Linux image on a flash drive using Windows. UNetbootin can be downloaded from *http://unetbootin.github.io/*.

The list of such utilities is never ending and they are increasing, day by day. One advantage of running Linux using a USB flash drive rather than a live CD is that you get

persistent data storage space, where some critical data can be saved and remains persistent.

Try Linux using a virtual machine

All the methods discussed till now provide a way to experience Linux, but if you really want to use Linux on Windows what you need is an environment which is persistent, more robust and gives you greater freedom to do more than just running Linux. Installing a Linux operating system as a virtual machine (VM) is the answer to this. VMs allow you to have multiple full-fledged Linux distros running on top of Windows, simultaneously. The only limitation to the number of Linux distros is your system's hardware resources. More importantly, all Linux operating systems will be completely isolated from the host Windows system—and from each other as well.

There are various ways to create VMs in Windows but the most common and free open source tool is VirtualBox. You can just download it and use it to create as many virtual machines as you want, to experience running multiple Linux distros simultaneously.

Try coLinux based distributions

Another way to run Linux on Windows is to use coLinux distributions. coLinux or Cooperative Linux is a very new way of running Linux on Windows natively, without any extra partition. Conceptually, coLinux is a port of the Linux kernel that allows it to run cooperatively alongside other OSs. With this, you will not require any virtualisation software to run Linux on Windows. Currently, it only allows the KNOPPIX Japanese version on Windows. But this is a project which is hot and offers a host of development opportunities to all open source enthusiasts.

More details on coLinux can be found at *http://www.colinux.org/* END

References

[1] *http://www.google.com*
[2] *http://www.howtogeek.com/170870/5-ways-to-run-linux-software-on-windows/*
[3] *http://www.labnol.org/software/run-linux-with-indows/19746/*
[4] *http://www.techhive.com/article/249870/create_a_bootable_linux_flash_drive_in_three_easy_steps.html*
[5] *http://www.ubuntu.com/download/desktop/create-a-usb-stick-on-windows*
[6] *http://www.2p.cz/files/2p.cz/downloads/howto/colinux.pdf*
[7] *http://colinux.wikia.com/wiki/Getting_Started_with_coLinux*
[8] *https://help.ubuntu.com/community/Wubi*

By: Ashish Bunkar

The author has over five years' experience as a Linux professional. His areas of interest are SAN and virtualisation technologies.

Going Live with FFMPEG

FFMPEG is a free and open source project that produces libraries and programs for handling multimedia files and data. This article shows how one can upload a video file, screen cast, or Web cam output for *YouTube Live* using FFMPEG.

Sometimes, advanced technology doesn't work too well. Streaming over the Internet is an example of this. Our old bulky TV sets play live channels in real-time without any complications while a computer needs 'this and that' to stream a single video. Playing video files on a website was an ugly job for both the developer and the user till HTML5 entered the scene.

Although video and audio streaming over the Internet has been made simple, live streaming is still challenging. But that is not the case with YouTube. Uploading video for live streaming is an easy task and the user has to do nothing more than own a device that can play regular YouTube videos.

Let us look at how we, as free software enthusiasts, can broadcast on *YouTube Live*. Let's use the highly sophisticated, portable, yet simple tool called FFMPEG. The command to invoke FFMPEG can be used in other situations too—to upload videos to sites other than YouTube which support the same technology. I have taken YouTube as the example since it seems to be the most stable and popular service.

How does it work?

The basic steps involved in *YouTube Live* broadcasting are somewhat as follows.

1. You have or get a verified YouTube account, which is in *good standing*.
2. You get the public server URL and private stream name/key from your *YouTube Live* dashboard, since it's required to configure things like monetisation.
3. Using these two values, you upload your video to YouTube with the help of an *encoder.*
4. The stream is broadcast live and the world gets to watch you.

RTMP and encoders

RTMP (Real Time Messaging Protocol) is an Internet protocol that can be used to stream audio, video and data between a Flash player and a server. It was initially developed as a proprietary protocol by Macromedia. But Adobe, which owns Macromedia now, has released an incomplete version of the specification for public use.

RTMP is implemented over the following three stages:

- Live video encoder
- Live and on-demand media streaming server
- Live and on-demand media streaming client

There are free software servers and clients available to make use of RTMP. Thus you should be able to replace Adobe Flash Media Server with Red5 or FreeSWITCH, and Flash Player with Gnash or VLC. But in the case of YouTube, Google takes care of the server and client. As a YouTuber, all you have to bother about is the *encoder*.

As the broadcaster, you have to encode and upload the video to YouTube using a supported encoder. Google lists some supported encoders in its support page. There you can see an open source suggestion, *Open Broadcaster Software*. You may try it, but here we are taking up FFMPEG.

Note: The latest version of Firefox is needed to enjoy YouTube's HTML5 player. Chromium users are recommended to install or update the package *chromium-codecs-ffmpeg-extra*.

Why FFMPEG?

Because it is the simplest option! Most distros have FFMPEG preinstalled; and all you have to do is run a single line of command and you are live on YouTube!

Also, FFMPEG is highly portable (it runs on different operating systems) and powerful. It supports a lot of video formats and configurations. It can take live input from devices and push the output to remote locations. You can feel its speed even on an old computer.

It is, therefore, used as a backend of many other GUI applications, including converters and video editors. The most important thing about it is that it is a free software package!

Note: Ubuntu now ships *avconv* instead of FFMPEG.

Let's broadcast a video

Let us go through a step-by-step 'how to'. It includes both YouTube and FFMPEG.

First, you have to prepare the video file. Although

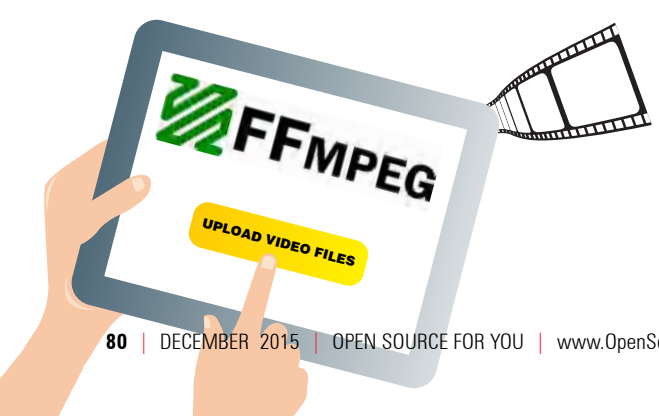

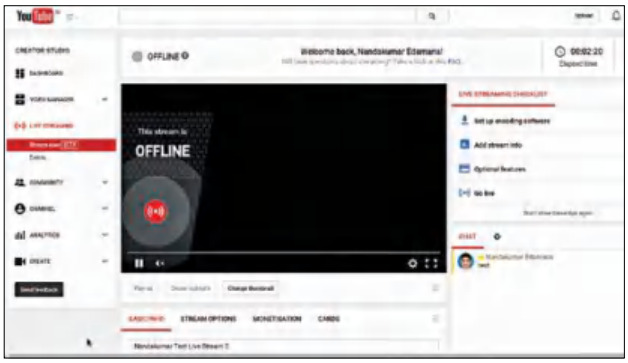

Figure 1: YouTube Live Dashboard

Figure 2: Encoder Setup – Get the Server URL and Key

FFMPEG can convert it for YouTube on-the-fly, it would be nice if you can render it for the YouTube format using a video editor or converter. This is because FFMPEG options are a bit messy and you don't want to go crazy at the uploading stage. Many video editors have YouTube presets (they also use FFMPEG as the backend!). For example, while rendering with OpenShot, take the profile as *Web* and target *YouTube-HD*.

Once you have your video file, sign into YouTube. You can now access the Live Streaming Dashboard using the link on the upload page or by visiting *https://www.youtube.com/live_dashboard*.

YouTube Live dashboard

Set things like the stream name and monetisation. At the bottom of the page, you can see ENCODER SETUP. Grab the server URL and stream name/key from there (by pressing the *Reveal* button).

Encoder setup: Get the server URL and key

Now open the terminal and use the following command:

```
ffmpeg -i <inputFilePath> -f flv SERVER_URL/KEY
```

Here is an example.
Server URL: *rtmp://a.rtmp.youtube.com/live2/*

```
Key: nandakumar96.7mwj-vh6r-45ru-6eqq
```

```
Command: ffmpeg -i /tmp/a.flv -f flv rtmp://a.rtmp.youtube.com/
live2/nandakumar96.7mwj-vh6r-45ru-6eqq
```

Your stream is online! You can now use a few options like shown below:

```
-s widthxheight -- to scale the video
-b value -- to set the bitrate
```

These two options are useful if you want to degrade the quality of the stream in case you have a low bandwidth. For example, you can stream an HD video file to 240p quality by using the following command:

```
ffmpeg -i /tmp/a.flv -f flv -s 426x240 -b 400k rtmp://...
```

Caution

Your live stream will be archived and uploaded to your channel once the broadcast is complete. To prevent this, you can edit the STREAM OPTIONS before starting the stream (in the live streaming dashboard).

Tip: A YouTube live stream is stopped when the encoder is terminated. You can press *Ctrl+C* to stop FFMPEG so that your stream also stops.

Make your screen public!

You've successfully broadcast a pre-recorded video. So why not try to screen cast it? Your screen is captured and broadcast live! The following command does that with your default microphone (or sound input device) as the audio source:

```
ffmpeg \
    -f alsa -ac 2 -i default \
    -f x11grab -framerate 24 -video_size 1366x760 \
    -i :0.0+0,0 -preset medium -r 24 -g 48 \
    -acodec libmp3lame -ar 44100 -threads 6 -qscale 3 -b:a
128k -bufsize 512k \
    -f flv -s 426x240 SERVER_URL/KEY
```

The above command is split into several lines for readability. You may enter this in a single line avoiding the backslashes (\). If you use them, make sure you don't enter white spaces before entering the next line. Also, don't forget to press the *Enter* key after completing the command.

video_size tells you how much of the screen should be recorded. *-s* tells you what should be the output video resolution. In the above case, my whole screen (1366x760) is captured and scaled to 426x240. If your connection is fast enough, you can omit the second option so that a high quality video is uploaded (there's no problem if the viewer has a slow connection — YouTube will take care of that).

So happy streaming! END

References

[1] *https://support.google.com/youtube/answer/2797387*
[2] *https://support.google.com/youtube/answer/2907883*

By: Nandakumar Edamana

The author is a writer, hacker and software freedom activist from Kerala. He has authored two books and developed many free software packages, which are available from nandakumar.co.in. His personal site is *nandakumaredamana.in*.

Using Open Source Software on Windows

Sometimes, for compelling reasons, one is forced to work on proprietary software like Microsoft Windows. Even then, one can take advantage of open source software designed for Windows. In this article, the author describes various software that can be installed on Windows, along with a few tips on how to use these effectively.

The site *www.netmarketshare.com* displays the current break-up of the desktop operating systems' market for July 2015 (Figure 1). At the outset, it's important to acknowledge the problems related to software, which are always the same. On my computer at home, I have the freedom to install my preferred GNU Linux distro, but if I work at a company, I'm forced to use Windows, which is so prevalent across enterprises. Generally speaking, at home, I can also use Windows, usually pre-installed, as a base on which it is possible to add free and open source software. The software shown in Table 1 can be classified, more or less, into the following three categories:

1. Free and open source
2. Free (or free for personal use only) but not open source
3. Commercial (only one: FaxTalk)

With the only exception of FaxTalk (in my opinion, nice and cheap), for personal use, the cost of using either of the above is equal to zero. This article will present some tips and tricks tested on Windows 7 and related to the software that is highlighted in bold in Table 1.

7-Zip	JRE • JDK
Acrobat Reader	Lazarus • NetBeans • Winpython
CDBurnerXP	LibreOffice
Emacs • jEdit • Notepad++ • Vim	Printer drivers
FaxTalk	**RedMon**
Firefox	Scanner driver
Flash Player	**Scilab**
GhostScript	ShellExView
GIMP	Silverlight
ImageMagick	Skype
Inkscape	System Utilities
Internet Security	**Thunderbird**
IrfanView	VLC

Table 1: Software for installation on Windows 7

Text editors

The importance of a text editor is well known. On Windows, a long list of text editors is available. This section describes how to build a syntax highlighting file for jEdit (*http://jedit.*

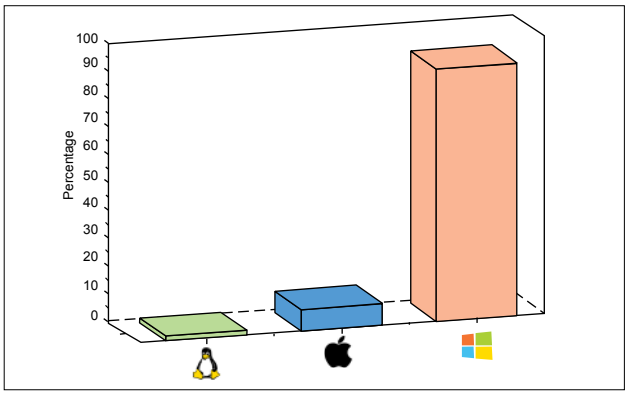

Figure 1: Desktop operating systems' market share

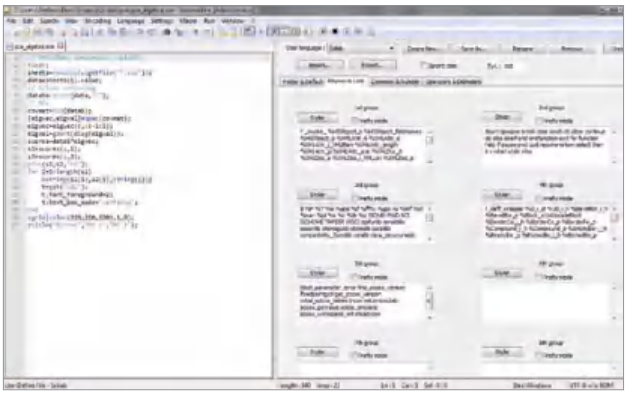

Figure 2: Notepad++ custom language editor

org) and for Notepad++ (*https://notepad-plus-plus.org*). Now, I would like to consider an imaginary programming language called *qwerty.* In the *qwerty* language, a comment line starts with a double slash (// this is a comment) and the strings are those between two quotation marks ("this is a string"). This language has many keywords which can be classified into various categories. In jEdit, the syntax highlighting file has the following structure:

```
<?xml version="1.0"?>
<!DOCTYPE MODE SYSTEM "xmode.dtd">
<!-- jEdit syntax file for the qwerty language -->
<MODE>
<RULES IGNORE_CASE="FALSE" HIGHLIGHT_DIGITS="FALSE">
<SPAN TYPE="LITERAL1" ESCAPE="\" NO_LINE_BREAK="FALSE">
<BEGIN>"</BEGIN>
<END>"</END>
</SPAN>
<EOL_SPAN TYPE="COMMENT1">//</EOL_SPAN>
<KEYWORDS>
<!-- Category 1 -->
<KEYWORD1>this-is-a-keyword</KEYWORD1>
...
<!-- Category 2 -->
<KEYWORD2>this-is-a-keyword</KEYWORD2>
...
<!-- Category 3 -->
<KEYWORD3>this-is-a-keyword</KEYWORD3>
</KEYWORDS>
</RULES>
</MODE>
```

Because I'm not interested in highlighting the digits, this option is defined as 'false'. Save the file as, for example, *qwerty.xml*, and put it in the *modes* directory (*C:\Program Files\jEdit\modes*). Then add the following line to the catalogue file located in the same directory:

```
<MODE NAME="qwerty" FILE="qwerty.xml" FILE_NAME_GLOB="*.qwe"
/>
```

For Scilab (*http://www.scilab.org*), the complete file about this first example is freely (as in 'freedom' and in 'beer') available at *https://github.com/astonfe/scilab.* In Notepad++, a syntax highlighting file is not very different from that for jEdit, and has a structure divided more or less into three sections. The comment line, math operators, punctuation, parentheses and some keywords, for example, in an if-else-end statement are in the first section:

```
<Keywords name="Comments">00//</Keywords>
<Keywords name="Operators1">+ - * / = &apos; , ; . : &lt;
&gt; ( ) [ ] { } | &amp; ~</Keywords>
<Keywords name="Folders in code2, open">for function if
select try while</Keywords>
<Keywords name="Folders in code2, middle">case catch else
elseif</Keywords>
<Keywords name="Folders in code2, close">end endfunction</
Keywords>
```

The keywords according to the following simple schema are defined in the second section:

```
<Keywords name="Keywords1">
this-is-a-keyword
...
</Keywords>
<Keywords name="Keywords2">
this-is-a-keyword
...
</Keywords>
<Keywords name="Keywords3">
this-is-a-keyword
...
</Keywords>
```

The strings, as well as the colours and fonts for each of the above categories, are defined in the last section. Save as *UserDefineLang.xml* and copy to *C:\Users\AppData\Roaming\Notepad++*. This editor also has its own built-in editor for custom languages, so it's not necessary to directly

edit (Figure 2). For Scilab, a complete file is written by Samuel Gougeon and is available from *https://fileexchange.scilab.org/toolboxes/140000*.

Firefox and Thunderbird

The user and not the software must determine whether or not to install updates. The use of internal options of each program may not work, so an alternative way is presented here. To disable the installation of add-ons, it is necessary to put two files in four directories, according to the following schema:

```
C:\Program Files\Mozilla Firefox\defaults\pref\local-settings.js
C:\Program Files\Mozilla Firefox\mozilla.cfg
C:\Program Files\Mozilla Thunderbird\defaults\pref\local-settings.js
C:\Program Files\Mozilla Thunderbird\mozilla.cfg
```

To turn off the updates of each application, one file must be added to your profile:

```
C:\Users\[your-username]\AppData\Roaming\Mozilla\Firefox\Profiles\[your-profile].default\user.js
C:\Users\[your-username]\AppData\Roaming\Thunderbird\Profiles\[your-profile].default\user.js
```

The contents of these three files are as follows:

```
// local-settings.js
pref("general.config.obscure_value",0);
pref("general.config.filename","mozilla.cfg");
// mozilla.cfg
lockPref("xpinstall.enabled",false);
// user.js
user_pref("app.update.enabled",false);
```

One of the easiest things about Firefox and Thunderbird is how to do a backup. I'm comfortable using the following batch script. A copy of each directory with the current date is created on the desktop.

```
@rem Backup for Firefox and Thunderbird
@rem Copy from source to destination including subdirs and hidden
@echo off
@for /F "tokens=1,2,3 delims=/ " %%A in ('date /t') do @(
set Day=%%A
set Month=%%B
set Year=%%C
set All=%%C-%%B-%%A
)
mkdir C:\Users\[your-username]\Desktop\Mozilla_%All%
C:\Windows\System32\xcopy.exe C:\Users\[your-username]\AppData\Roaming\Mozilla C:\Users\[your-username]\Desktop\Mozilla_%All% /S /E /H
mkdir C:\Users\[your-username]\Desktop\Thunderbird_%All%
C:\Windows\System32\xcopy.exe C:\Users\[your-username]\AppData\Roaming\Thunderbird C:\Users\[your-username]\Desktop\Thunderbird_%All% /S /E /H
echo Backup completed.
pause
```

GhostScript

To configure a PDF printer, first install GhostScript and RedMon, the Redirection Port Monitor. I have installed GhostScript in *C:* so its directory is *C:\gs9.15*. The next step is to add a new printer on the port RPT1. The driver for the HP ColorLaserJet 2800 PS works because it's both colour and postscript. For this printer, choose a name you like – for example, GS2PDF. Then put the following code in the file *C:\gs9.15\pdfwrite.txt:*

```
-IC:\gs9.15\lib;C:\gs9.15\fonts
-sDEVICE=pdfwrite
-dPDFSETTINGS=/prepress
-dCompatibilityLevel=1.4
-dNOPAUSE
-dSAFER
-sPAPERSIZE=a4
-r600
```

The last step is the configuration of the RPT1 port on which the GS2PDF printer is installed, as shown in Figure 3. Also, under *Device Settings*, set *Send Ctrl-D* before each job to *'No'* and *'Wait Timeout'* to 1 second. The arguments for GhostScript are as follows:

```
@C:\gs9.15\pdfwrite.txt -sOutputFile="%1.pdf" -c .setpdfwrite -f -
```

This first part about GhostScript is based on the notes written by Michael B. McClelland (Reference 1) and by Henrik Schmiediche (Reference 2). On PDF files, there are some common operations that are easy to carry out on Windows. If I need to extract some pages from a large PDF file, I use the following batch script. It's only necessary to specify the first and the last page numbers to extract the input and the output file names.

```
@rem PDF extract
@echo off
cd C:\gs9.15\bin
gswin32c.exe -sDEVICE=pdfwrite -dPDFSETTINGS=/prepress
-dNOPAUSE -dBATCH -dSAFER -dQUIET ^
            -dFirstPage=2 ^
            -dLastPage=6 ^
            -sOutputFile="C:\Users\[your-username]\extract4.pdf" ^
```

```
            "C:\Users\[your-username]\input.pdf"
echo Done.
pause
```

Another common operation is the merger of some PDF files to obtain one single file. In this case, I use the following batch script, in which it's necessary to specify the output file name and the file names of each file to merge them.

```
@rem PDF merge
@rem Input as single files "extract1.pdf extract2.pdf ..."
@echo off
cd C:\gs9.15\bin
gswin32c.exe -sDEVICE=pdfwrite -dPDFSETTINGS=/prepress
-dNOPAUSE -dBATCH -dSAFER -dQUIET ^
            -sOutputFile="C:\Users\[your-username]\output.
pdf" ^
            "C:\Users\[your-username]\extract1.pdf" ^
            "C:\Users\[your-username]\extract2.pdf" ^
            "C:\Users\[your-username]\extract3.pdf" ^
            "C:\Users\[your-username]\extract4.pdf"
echo Done.
pause
```

There is also another operation that's sometimes useful - the conversion from a PostScript file to a PDF file. Even in this case, GhostScript can be the solution. In the following script, it's only necessary to specify the file names of the input and output:

```
@rem PS2PDF
@echo off
cd C:\gs9.15\bin
gswin32c.exe -o "C:\Users\[your-username]\output.pdf" ^
            -sDEVICE=pdfwrite -dPDFSETTINGS=/prepress ^
            -dHaveTrueTypes=true -dEmbedAllFonts=true
-dSubsetFonts=false ^
            -c ".setpdfwrite <</NeverEmbed [ ]>>
setdistillerparams" ^
            -f "C:\Users\[your-username]\input.ps"
echo Done.
pause
```

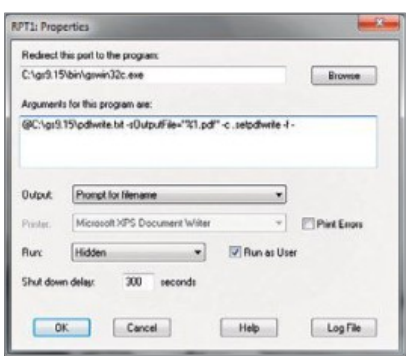

Figure 3: Arguments for GhostScript

GNU Emacs and ImageMagick

There is a lot of software that can be used as a personal finance manager. Some examples are: GnuCash, HomeBank, KMyMoney and Skrooge. I have tested them on GNU Linux and not on Windows, so I can't say more about their installation and use on Windows. This kind of program is nice but I'm not satisfied with it.

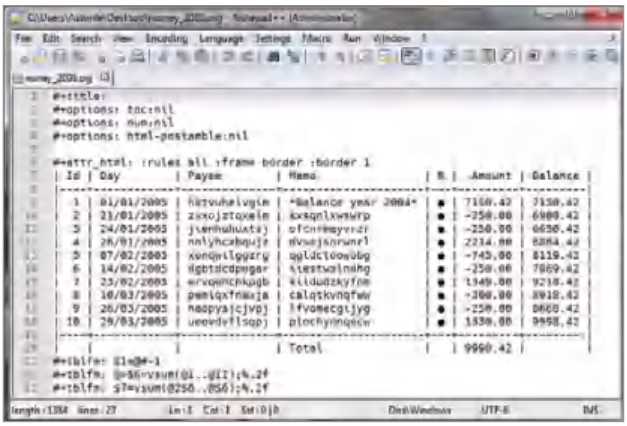

Figure 4: An *org-mode* file opened in Notepad++

Another option is to use something like Ledger (*http://www.ledger-cli.org*), which is available for Windows too. But I prefer to follow a personal method and use *org-mode* to build some tables. With these tables, I can't do queries and reports, but I don't need them. A typical table I use has seven columns:

1. ID, which is a progressive integer
2. Day, which is the date when the operation is done
3. Payee, which is the payee
4. Memo, which is the explanation of that operation
5. R, which is used to show if the operation has been reconciled or not
6. Amount, which is the amount in that operation that can be a positive or a negative number
7. Balance, which is the running total

The first row is the total balance for the previous year. At the end of the table, there are the following three rows:

```
#+tblfm: $1=@#-1
#+tblfm: @>$6=vsum(@I..@II);%.2f
#+tblfm: $7=vsum(@2$6..@$6);%.2f
```

The first adds the ID number (column 1), the second calculates the total balance value (the sum of all the values in column 6) and the third calculates the running total (column 7). This is a simple, practical and very fast way to open a file on a text editor and to print as *html*, when necessary. Figure 4 shows an *org-mode* file opened in Notepad++ with the free and open source font Hack v2.010 (*http://sourcefoundry.org/hack*). Regarding ImageMagick - after the installation, to verify if it's working properly, type each of the following lines at the command prompt:

```
convert logo: logo.gif
identify logo.gif
imdisplay logo.gif
```

If the file *mfc100u.dll* is missing, it's impossible to run *imdisplay.* Simply download and install Microsoft Visual C++ 2010 SP1 Redistributable Package. Some useful common operations are summarised in the following code from the command prompt. ImageMagick is very powerful and has many capabilities. More examples can be found at *http://www.imagemagick.org/Usage*

```
- Reduce image quality/size
  convert -strip -quality 25 big_image.jpg new_image.jpg
- Generate a canvas with a specific size and color
  convert -size 1280x800 canvas:"#708090" background.jpg
- Convert from PNG to ICO
  convert -define icon:auto-resize=256,128,64,48,32,16 file.
png file.ico
- Convert from PDF or EPS to PNG
  convert -density 300 -quality 100 -layers flatten file.pdf
file.png
- Convert from color to gray scale
  convert -colorspace Gray color_image.png gray_image.png
- Modulate Brightness, Saturation, Hue
  convert -modulate 100,100,-50 image.png new_image.png
- Resize
  mogrify -resize 128x256 image.png
```

Before the final conclusions, I would like to add some remarks about these three kinds of software: database, CAD and RAD. For a database on a network, I prefer to use MySQL/MariaDB with some PHP (see the *OSFY* August 2015 issue) and for Drupal websites too. But for a small, desktop only application the story is different. I like well-known commercial software, but find them expensive. But one well-known commercial software that is not too costly is LibreOffice Base. It is an interesting but, in my opinion, not so good alternative to the above-mentioned commercial software, because it still needs a lot of work before it can be considered a true alternative. The Java engine is slow; Firebird has some stability problems and the GUI is not the best. So, actually for me, this is an open topic.

About CAD, I have no experience about 3D so I would like to list six interesting free alternatives to AutoCAD LT and to VectorWorks (I haven't used them both for years): DoubleCAD XT, DraftSight, FreeCAD, Medusa, nanoCAD and Solid Edge 2D. Some of these options also have 3D capabilities but only FreeCAD is open source. I've tested only Solid Edge (because it's used in the CAD/CAE course at my college) and DraftSight, and I think that the second, for my simple needs, can be an alternative because it's available for GNU Linux (note that DraftSight is unsupported on Linux 32-bit machines).

Last, some short notes about NetBeans and Lazarus. As well as in the case of text editors, I can't make a definitive choice because Java is very widely used and having some knowledge about it has certain advantages, but I also like the Pascal syntax. At this point a big section could be open to talk about the Python programming language on Windows. But this is another story. Finally, can Windows be a good base for free and open source software? I think it can, because on Windows, I can install many other free software (even if not open source) or, if necessary, commercial software that's not available for GNU Linux. But for the future, I hope that it will be possible to use only GNU Linux. END

References

[1] *http://mbmcclelland.blogspot.it/2012/11/creating-free-pdf-printer-in-windows-7.html, last visited on 28/08/2015.*
[2] *http://www.stat.tamu.edu/~henrik/GSWriter/GSWriter.html, last visited on 28/08/2015.*

By: Stefano Rizzi

The author works in the analytical chemistry and textile chemistry areas. He has been a Linux user since 1998.

Daily Use FREE SOFTWARE for Windows

For many Windows users, the lure of pirated software is often difficult to resist. But there is no need to break the law when legal, open source and free software that is sometimes even better than the package it is replacing, is available. This article lists out free software that can be used on a Windows machine to replace expensive proprietary software.

Let's look at some essential software that is either free or open source and can replace many of the expensive software packages on Windows. Not only are these packages free but they are also easy to use. Almost all open source packages are available in the form of executables or exe formats, to make installation easy on the Windows platform. One can always have a look at the source code, which can be downloaded from the respective website.

Typically, some software come preinstalled with a Windows PC or laptop, while the rest are installed as per individual needs. The most commonly used software are Internet Explorer, Microsoft Office (which includes Word, Excel, Powerpoint, Outlook, Publisher, Access, Media Player for playing back audio and video files); compressing software like Winzip, Winrar; different tools for instant messaging and video chat, like Messenger provided by Facebook, Yahoo, etc; Skype for video chat and Adobe Acrobat for PDF reading. Some users also like to install Photoshop, Dreamweaver, Illustrator, CorelDraw, Maya, etc, for designing and animation purposes. These are a few commonly used packages and almost all are proprietary – one needs to buy them to use them. The cost is huge and Microsoft Office alone can cost about Rs 15,000 for the professional version. Some other software are free but come with adware or spyware.

The following section lists some of the free and easy-to-use alternatives to paid software.

Web browsers: Firefox and Chromium are the front runners in replacing Microsoft's Internet Explorer. Both are free and open source, and come with many add-ons which make surfing the Web better and smoother. Firefox, in fact, is the default browser of many well-known Linux distros.

Email clients: Thunderbird and Evolution are well known, being the default email clients in many Linux distros. Thunderbird is easy to configure, easy to use and lightweight. Like Firefox, it also comes with many add-ons, which include spam filters, protection from phishing attacks and a lot more. It can replace Microsoft Outlook.

Communication and chatting: Pidgin and Empathy are instant messaging clients. They can replace Facebook Messenger, GTalk, MSN Messenger, Yahoo Messenger, etc. Xchat is used as an IRC client, and is also free and lightweight compared to its proprietary counterparts. Skype, a popular Voice-over-IP, can be replaced by Ekiga (Gnome Meeting), which supports HD quality sound and video.

Productivity: The most expensive packages come under this category, e.g., Adobe products, the Office suite, desktop publishing software, etc. There are many open source software packages that can substitute these expensive Windows packages. OpenOffice and LibreOffice can replace the Microsoft Office package which includes Word, Excel and Powerpoint. Apart from this, Abiword is a replacement for Microsoft Word; Scribus replaces Microsoft Publisher; PDF Creator replaces Adobe Acrobat; while FengOffice provides project management, client

management, task management, workflow processes, time tracking, document management, reports, a calendar and many other features. It is available as a community version (free) and a paid version.

Drawing: GNU Image Manipulation Program (the GIMP) can do all kinds of image manipulation and is the best alternative to Adobe Photoshop. Dia, a diagram editor, can be used for drawing diagrams, as an alternative to Microsoft Visio. And Inkscape easily replaces Adobe Illustrator and CorelDraw.

Science: SciLab is the best alternative for MatLab.

Antivirus: ClamWin is an antivirus software with a slick design and is easy to manage.

Calendar: For those who need calendar support, Sunbird provides an excellent calendar for Linux distros. It is easy to use and master.

Animation: Blender is great for 3D animation and is the best replacement for the Maya package. There are many movies that have been made or rendered using Blender.

Audio and video: Playing audio and video on the Windows platform means installing many software packages for different formats. VLC is the solution for all requirements. It supports many audio and video formats and can play back almost all audio and video files. It is simple to use and has a minimalist design. It can also be used for various streaming protocols and for media conversion. Audacity is another tool used for audio editing, as a replacement for Adobe Audition.

Utilities: For torrent downloads, one should try the lightweight BitTorrent client. For compressing, 7-Zip is an excellent option. For FTP, the Filezilla client is a good option.

Database: All major database software like MySQL or PostgreSQL come free as community editions, and can replace the costly Microsoft SQL Server and Oracle Server.

Virtualisation: Virtual Box is desktop virtualisation software that can be used in place of VMware.

So, basically, for every proprietary software, there is an open source software option, which is usually distributed free as a community edition, but may have a paid version too. It's simply a matter of searching for what you want. END

By: Ashish Singh Bhatia

The author is a technology enthusiast and a FOSS fan. He loves to explore new technology, and work on Python and Java languages. He can be reached at *ast.bhatia@gmail.com* and blogs at *https://openfreeidea.wordpress.com/*

Building Effective Machine Learning Solutions with Open Source Tools

Extracting hidden insights and patterns from data is of immense value to any business. Machine learning addresses this need effectively. But for it to work well, the right tools need to be used. Let's take a look at a few open source tools that could be put to good use in this domain.

The adoption of machine learning is not only due to the availability of more cost-effective and powerful data acquisition technologies and processing capabilities in the hardware domain. The proliferation of many tools, especially in the open source domain, makes machine learning easier to implement.

Along with the right set of enablers, machine learning is becoming pervasive in both enterprise as well as consumer applications. The usage of this technology can really be widespread, ranging from telematics to product recommendations, and from speech recognition to condition monitoring. No wonder, machine learning and the related technologies are rapidly becoming an integral part of the enterprise technology stack. Today it is considered as a critical skill of the developers as well as of the architects.

While machine learning can act as a key enabler to any business, for it to work effectively, it is crucial to choose the right set of tools. There are a wide range of tools available in the open source domain. Each one of them has evolved to cater to some specific need and has its own set of use cases. It is not fair to do an apple-to-apple comparison of all tools available today. Rather, it is necessary to analyse and understand the business need at hand and, accordingly, choose the tool that would be the best fit to solve the problem.

In this article, with the help of some relevant examples, we will discuss some of the key machine learning use cases along with the open source tool sets suitable for each of them.

Machine learning: A quick preview

Before delving into the machine learning tool set, let us take a quick look at what machine learning is all about and how it can be useful for business. Machine learning involves a system that acquires intelligence from past data, experience, observation and training. As part of machine learning, statistical techniques are used to learn the patterns, trends and structure hidden in the data. The learnings or the insights are then leveraged to perform many business operations. For example, by identifying books that have been purchased by your friends, Amazon is able to recommend books to you. Similarly, if you have a habit of watching the latest action movies on Netflix, it will suggest other action movies for you, by using your past viewing history.

Open source tools

Till recently, machine learning was confined more to the world of academics. That may be the reason for the plethora of open source tools and libraries around machine learning and related technologies. With the passage of time and the vast interest in this area, many open source tools and libraries have evolved. By simply using the most suitable implementation of machine learning algorithms, it is possible to derive insights fairly quickly, easily and in a cost-effective manner. Some of the popular tools are listed below.

- **R**: This is a language and environment for statistical computing and graphics. R provides a wide variety of statistical (linear and non-linear regression, classical statistical tests, time-series analysis, classification, clustering, etc) and graphical techniques, and is highly extensible. The R language is often the vehicle of choice for research in statistical methodology, and it provides an open source route to participation in that activity.

- **Python**: While R is specifically created for statistical analysis, Python also has a rich set of machine learning implementations. It is widely used among the scientist community. Being an interpreter, high level programming language, Python is a good fit for machine learning implementation, as quite often, these implementations call for an agile and iterative approach.
- **Apache Mahout**: This provides an environment for quickly creating scalable, performant machine learning applications.
- **H2O**: This is for data scientists and application developers who need fast, in-memory scalable machine learning for smarter applications. H2O is an open source parallel processing engine for machine learning.
- **RapidMiner**: This is a platform that provides an end-to-end development environment for machine learning. Through a wizard-driven approach, RapidMiner allows the user to quickly build the predictive analytics model.

Choosing the right tool: A key challenge

It is extremely crucial to choose the appropriate tool for the successful implementation of machine learning tool. As seen in the previous sections, there are a wide range of open source machine learning tools available. Each of these tools has its own feature set, strengths and limitations. Due to the varied feature sets, it is not possible to directly compare the tools or suggest one that can be considered as the best fit for all use cases. Hence, it is often a challenging task to select the right tools set to solve the business problem at hand. We need to first have a clear understanding of the problem. Based on that understanding, we need to choose the right tool set which is the best fit to solve that type of business problem.

Like in any other domain of software architecture, selecting a machine learning tool should be driven by the business use case. To choose the right tool, we first need to understand the business context and the business driver. Once we define the problem and get the consensus from all the stakeholders, we need to identify the machine learning technique that would be suitable to meet that business driver. The identified functional and non-functional requirements would form the basis of choosing the appropriate tool which would cater to the business need.

Now, if we consider the entire spectrum of business use cases around machine learning, *prima facie* they can be grouped into three segments:

- The classical use cases, which call for in-depth data analysis and very specific algorithms. In these cases we need to use tools like R, Python and Octave as these tools have a huge set of implementations for a really large number of algorithms. These tools allow us to create a bespoke solution by selecting a specific algorithm. Also, with the basic building blocks given in these tools, the user often customises the algorithm or the implementations to address a specific business need. As these tools have evolved mainly due to the work of and for the data scientist community, most of the time, they may not have a very user-friendly interface and the learning curve may be relatively steeper.
- The second set of use cases have a strong focus on scalability and performance, along with unveiling hidden insights. This is popularly known as big data analytics. A default choice for these use cases would be suitable tools from the Hadoop ecosystem. A widely used tool in this category is Mahout. It is one of the pioneers in this league. Recently, we have been seeing many more improved and optimised tools in this category like H20, Sparkling Water, MLib, etc. These tools leverage distributed computing and in-memory technology to come up with faster and scalable implementations of machine learning algorithms. The choice of algorithm is much less in these tools. Also, quite often, for the desired results, they require high-end infrastructure.
- The third segment of use cases is primarily for the business users' community. Here, users want to get some quick insights not only by visualisation but also from the predictive analytics perspective without really getting into the details of the statistical techniques or algorithms. Of late, we are seeing a lot of development in this domain, with many vendors coming up with their offerings — Microsoft Azure, Amazon machine learning, IBM Watson Analytics, etc. Many of the tools in this space are in the commercial domain. The open source tools like RapidMiner and Weka can be extended, and would be a good fit to satisfy this requirement.

Moving ahead

To trigger wider adoption of machine learning, many of the tools are now being deployed on the cloud. These cloud based tools come with their own pros and cons. They reduce the concept-to-cash window even further, as with these tools, the user does not have to worry about the hardware and the deployment aspect. Here, the user can concentrate more on the business solution and can utilise resources more effectively. Today, most cloud based tools are available in the commercial domain. Some open source tools like RapidMiner are available as cloud based deployments. These tools would be a very good fit for the third category of use cases. It is a good idea to watch the developments in this domain.

With the wider adoption of machine learning, businesses are expecting value from their investments. For effective implementation and to obtain the expected return on investment, the right choice of tools is very crucial. END

By: Sanghamitra Mitra

The author is a senior technical architect with more than 15 years of experience in the IT industry. She has worked with international clients like British Airways and Hitachi Medical Corporation. Her current focus area is solving business problems with machine learning and Big Data analytics.

www.indiaelectronicsweek.com

An EFY GROUP EVENT

SHOWCASING LATEST T&M PRODUCTS & LEADING T&M BRANDS

Attend India's First LIVE Demonstration of Latest T&M Equipment

TEST & Measurement India

January 11th-13th, 2016 • Bengaluru

To register, visit: **register.efyexpo.com** or call 011-26810601/02/03

To increase your sales by exhibiting in this event, Call: **Jyotsna (011-40596605)** or write to us at: **efyexpo@efy.in**

Co-located at INDIA ELECTRONICS WEEK | January 11th-13th, 2016 Bengaluru

OpenStack Private Cloud Adoption: The Value-added Systems Integrator's Role

The cloud is not a novelty any more and the IT world is more or less familiar with its potential and pitfalls. But the new entrant to the cloud family, OpenStack, has disrupted the whole ecosystem and is making headlines every day.

The open source legacy of the OpenStack building blocks and the availability of a wide variety of OpenStack 'builds' – from the free community distributions to strictly adhered to Service Level Agreement associated versions – have definitely boosted the implementation of cloud based computing infrastructure, especially in the enterprise context.

As a matter of fact, OpenStack is basically providing a set of application programming interfaces (APIs) to orchestrate the infrastructure configured under its disposal-say, compute, network, storage, security, etc – which will help to set up a cloud infrastructure resembling the characteristics of all major and widely known Infrastructure as a Service (IaaS) cloud platforms in a more efficient and economical way.

The word 'cloud' instantly brings up a lot of questions and concerns about its security and availability aspects.

Once those concerns are addressed and the readiness of the cloud based IaaS model technically accomplished, the next logical step is to identify the applications to be boarded on to the cloud and its eventual scalability as business grows or during the computing peak time.

Systems integrators, primarily those who specialise in cloud solutions or have a dedicated delivery model around cloud integration and provide value added systems integration services, are playing a key role in the adoption of the OpenStack based private cloud. This will help enterprises to address many of its infrastructure deficiencies.

Value added cloud integrators should always have an applications based approach rather than an infrastructure-centric approach. A bench of computing infrastructure, orchestrated via OpenStack is mostly an academic exhibit, without a set of cloud style applications deployed on it and serving the business in an efficient fashion in terms of nullifying many of its earlier associated issues like agility, elasticity, metering, scale, etc.

The cloud integrator can provide a quickly deployable appliance architecture approach to deploy a test and development private cloud for enterprises to fast-track their 'cloud-enablement' of the applications and execute proof-of-concept roll outs to convince its internal stakeholders and also to launch a pilot production environment.

Another major value addition that cloud integrators can bring in is a 'Suite of Tools' as a service, and its easy execution model to analyse and report the security compliance and performance characteristics of the implemented cloud infrastructure in the respective domain context (e.g., PCI compliance of the proposed design to a financial organisation or to an e-retailer).

Even though a private cloud brings in more elasticity and a granular metered view to the internal computing infrastructure usage, it is always finite in terms of its underlying data centre specifications and availability. This may defeat the very purpose of scalability, which enterprises are expecting out of a cloud enabled infrastructure investment.

Value added systems integrators can present a hybrid approach to the private cloud architecture, to rent its spikes to an external infrastructure service entity without hampering the compliance guidelines around which the private cloud is built on. This is a major boost to justifying cloud adoption and its added financial overheads against the end goal of the return on investment (RoI).

But as a bottom line, even though the aforementioned cloud implementation approaches and services are supposed to be the prime components of a value-added cloud integration proposition, the financial model and flexibility it offers to share the risk of a cloud adoption decision (mostly due to the nature of the application ecosystem and tight investment guidelines) of the enterprise, should be considered while choosing the cloud integration partner. END

By: Anuj Bhalla

The author is the VP and global business head for systems integration, maintenance services, and products in global infrastructure services, Wipro Limited. His recent initiatives have been the Wipro Open Data Center and the SDI Center of Excellence launched recently at the Bangalore Wipro Campus. The Center of Excellence is a first for any systems integrator or technology provider in the country.

"We get a lot of good talent when it comes to open source"

Ask a CTO about what the biggest problem is when it comes to using and implementing open source, and the answer would be 'manpower crunch'. But ***Piyush Jha, AVP, product engineering, GlobalLogic,*** has a different opinion. He finds enough talent to carry out his open source projects. In fact, he thinks that the modern day techies vouch for open source and the power they get with it. ***Diksha P. Gupta*** from ***Open Source For You*** spoke to Jha about his open source implementations, while inviting him to ***Open Source India 2015.*** Excerpts:

Q Do you use open source technologies at GlobalLogic?

A lot! We develop a lot of components over open source. My background has been mostly in Java. When we try and decide on the technology for a project, we often go with open source since that's a pretty comfortable zone for us. Apart from Java, we also work with open source languages like Ruby, Python, etc.

Q When you start developing a product, what's your first choice?

We work with a lot of partners and develop products for them. Being product developers, we need to get into the ecosystem of our partners, and try and see what they are doing. So, we generally try and propose a tech stack that mostly matches with the tech stack of the partner. Let's say a company is a complete .NET enterprise, then we have to propose technologies that are closer to .NET. We try and propose a tech stack that matches theirs. But if the decision is really open, then we carry out research on which technology suits their needs better.

Q When it comes to your clients, how much of acceptance is there for open source?

There is a lot of acceptance. What the clients want is that the service provider or product development partner should be well versed in the particular technology that's being recommended and be able to back that up fully. So, if we show them that we do have enough expertise in what we are

talking about, then things get easier. Open source won't get the desired boost if we do not have the backing of enterprise-led activation. So, if we have some successful case studies with respect to open source, then convincing clients is not a challenge.

Q How has your journey with open source been so far?

It's been extremely good. Our engineers are pretty comfortable in approaching forums, looking for a solution, trying to resolve a problem that has not been resolved before and getting their problems solved. People go for closed source solutions or tried-and-tested solutions if they are not confident about finding a way out of a problem. Our team is pretty good at going to forums to seek solutions, and it tries to find out how to do what others have not done before.

Q What are the pre-requisites that one should have before going for an open source solution?

In general, you have to break the problem into parts. So, even if the specific product that you are attempting to implement is something that you haven't done before, you would like to try and see whether its components or pieces have been already worked on before or not. If the pieces have already been worked on, it is a cakewalk. But if they haven't been worked on, then you would pick up a couple of pieces that you find slightly challenging or those that slightly deviate from normal coding and try and do a small proof-of-concept (PoC) around it before adapting it as the solution for the complete team or the clients to follow. So, these PoCs are pretty relevant in the sense that they take away a lot of the risks.

As an example, we are currently working with a leading retailer in India to try and automate its warehouse management system. At present, its systems are mostly on paper or Excel spreadsheet – that is how the picking, packing, shipping process is carried out. We are in talks with the firm to try and automate the process. When we proposed the solution, we found a couple of elements that they are not really confident about and which they haven't worked on before. And those are the parts of the solution that were pretty exciting to them. We took a step back and told them that these are parts in the complete solution that we haven't really done before so let us get a proof-of-concept on these points. So, we are currently doing a one-month project rather than a complete one. Once we are done and agreed on what the thresholds and boundaries would be, we will go for the complete solution.

Q What are the advantages of using open source technology?

Cost is one definite advantage. Apart from that, we see that generally, updates come pretty fast in open source as compared to a company that would be developing a particular product. There's an army of people who support it and keep contributing to it, which makes it pretty easy for any problem or any new idea that comes in to get implemented in the open source project. So, updates, adaptations or innovations come out first in the open source world. Whatever is coming in as a standard solution gets adapted pretty fast but besides that, there are quite a few innovations that come out and we get the first-mover advantage when we roll out innovations that people haven't seen.

Q According to you, what are the disadvantages of using open source technology?

You don't have solutions readily available all the time. At times, the support is limited. And you could get into some tight corners which are difficult to come out of. So, it's not easy to get a vendor from that corner for support. It would be slightly complicated since you have to work your way into forums, try and get your queries answered, etc. Nobody might have worked on the solution you have in mind so you might need to work around the problem.

> "It's easy to get people. Also, I believe, that the number of people in open source is either in equivalent to or higher than those using proprietary technologies."

Q Do you get enough talent to work on open source based projects?

The open source projects are getting increasingly exciting. The techies and even those who've just passed out of college are excited about open source and are willing to try their hands at it. Moreover, it's not really tough to motivate people to work with open source.

Q Are these people talented and skilled enough?

Yes. Because the geeks are more excited about working on open source, we get a lot of good talent when getting started on open source projects. Interestingly, they keep themselves updated on what is happening in the world and would know a lot of stuff that our clients may not have heard of. So, it's easy to get people. Also, I believe that the number of people in open source is either equal to or higher than those using proprietary technologies.

Q You are one of the few CTOs to say so. Most people feel that talent is one of their biggest problems...

We are seeing a lot of traction in the Java world today. There are a lot of firms that are working more on Java currently. If we look at Indian e-commerce players, all of

them are Java sourced, like Flipkart, Snapdeal or Myntra and if you look at their tech stacks, it is a lot of open source. You might think I am one of the few telling that talent is easy to find, but we do have quite a few companies who are actually able to find and use this talent.

Q When did your tryst with open source technology begin?
We started to do small projects during my third year of college in IIT. We did a lot of niche PoCs and then I did my internship at the University of Paris. There were a couple of professors there who were extremely pro-Java and pro-open source. So, working with them for three months actually exposed me to a lot that I was unaware of earlier. That experience got me more interested. Over the first four years of my job, I worked on five or six different languages including C++, Java, Perl and Ruby.

Q What's your primary language right now?
None. I am supposed to be a pro, across the board. To me, there are a lot of solutions that themselves shout out for certain languages to be implemented. So, if you want to build an extremely scalable enterprise solution that supports 10,000 concurrent users, or a million users or hits at a time, per day, you would probably go with Java. However, if you want to go to market pretty quickly with a website, and your investment pressures or your go-to-market strategy forces you to get that done within 1.5 months, you will probably not go with Java and select Ruby instead. So, it really depends on what your product vision is and what's the technology that your product is looking for. The decision is taken on that basis. Basically, the technology choice, more often than not, lies in the problem statement itself.

Q How easy or difficult is it for you to convince the management about using open source?
We are 50 per cent open source and 50 per cent proprietary. So, that has never been an issue. The management has always been convinced about it. Good techies guys have founded this company and have been around for quite some time, so open source has always been in the company's blood. Even today, our management is pretty supportive of what the technology team says. When members of my engineering team back up the solution and say that they can do it, we rarely face any resistance from the management.

Q Are there any community contributions that your people do?
Yes. A lot of our people are pretty active on blogs and forums. At times, the CTO office is involved in global capacity contribution but at an individual level too, we contribute pretty often. In GlobalLogic, senior people have a part of their salary as a bonus component. More often than not, 15-20 per cent of that component is earned by their contribution to the blogging world, the open source world and the world of technology. So, we have a concept wherein a part of the bonus is reserved for performances or contributions that they have made to the project, to the company or to the technology world, in general.

Q Twelve years back when you started, was open source technology as good as it is today? What were the challenges then?
I think it was even better than today. Now, there are a lot of proprietary solutions available and the world has evolved a lot. Even Java solutions today are proprietary. At that time, Sun was completely backed by open source and, today, I would say that the market is 25-30 per cent open source. At that time, it was at least 50 per cent open source. There was a lull that we saw in the middle. But in the past two years, with the mobile revolution coming in, the open source wave has started to pick up momentum again.

Q How much do you see it changing in another 3-5 years?
I think it's in a stable state. It is here to stay because the IT world today is completely oriented around innovation. There's a lot of competition around products. Let's take the example of e-commerce products. There are at least 10 portals in India, of which three or four are in the top league and they are pretty close in terms of their competitiveness. Their only differentiation is in the kind of features and the ease of use they provide.

Any company amongst these that comes up with two or three attractive features first, gets the advantage of taking a leap forward. And in such sectors, whenever you mention innovation, open source is kind of a front-runner, by default, . There are a lot of people and good thinkers around who think up something, write code for it and put it in. So, whenever you talk about innovation, a lot that's new and path-breaking would be open source. So, I think in the next two-three years, we will see a lot of evolution in open source.

Q What are the areas in which you see open source growing a lot?
If you talk about domains, retail e-commerce, mobility and media will continue to have a lot of open source elements. IoT is an extremely vast playground, where open source will play a very large role because that's a field which is still evolving. There will be a lot of media solutions. Media is changing a lot. There are very few people who read news on a desktop today. It's all mobile. People read the news, like, share, cut elements, etc, all on-the-go; so this is another area that will have a lot of open source contribution.

Q Is open source still a back-end property?
I really don't see any reason why that would be the case. In fact, many of the front-end operations use mostly open source, because it's much more feasible to do so. END

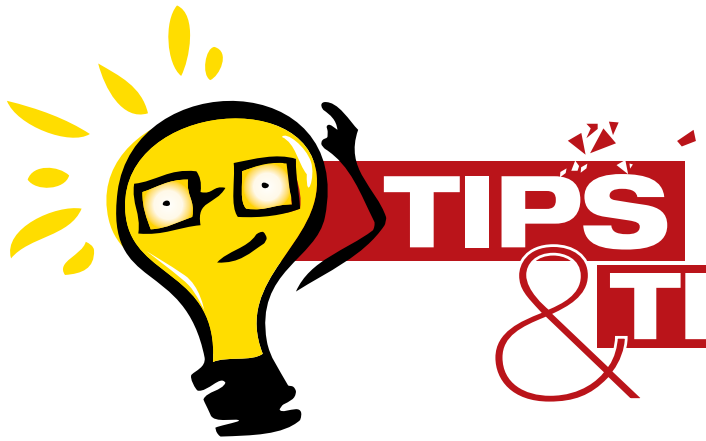

How to show mounted filesystems

If your *df* output is not showing all mounted filesystems, then here is what you can do. The *df* command pulls the information from the file */etc/mtab*. Sometimes the *mtab* file is corrupted, missing or has incorrect entries. So the *df* command output shows up incorrectly. However, the file */proc/mounts* keeps track of all the *mountpoints* correctly.

There are two workarounds for this issue.

a) Reboot the server and the file */etc/mtab* will be recreated with the correct entries.

b) If rebooting is not an option, then create symlink */proc/mounts* to */etc/mtab*.

```
ln -sf /proc/mounts /etc/mtab
```

—Krishna Murthy Thimmaiah, krishnamurthy.gt@gmail.com

A Python program for the recovery of deleted files from Trash

Here is a small Python code snippet that can recover files from Trash:

```
import os
import re

path="/home/gurukul/.local/share/Trash/files"
infopath="/home/gurukul/.local/share/Trash/info"
dirlist=os.listdir(path)   #list of file which present in file
folder
directory=[]
popis=""
for fname in dirlist:
    directory.append(fname)
    popis=popis + " " + fname
print popis
fname=raw_input("\nEnter the file name which toyou want to
recover")

a=open(infopath+"/"+fname+".trashinfo","r")

for line in a:
    if "Path=" in line:
      ab=re.findall(r'/.*',line)
destipath=str(ab)
destipath= destipath.lstrip('[')
destipath=destipath.rstrip(']')
destipath=destipath[:-1]
destipath=destipath[1:]
print "destination path is"+ destipath
file1 = open(path+"/"+fname,"r")
file2 = open(destipath,"w")
file2.write(file1.read())
file1.close()
file2.close()
print "files is recovered to desination"
os.remove(path+"/"+fname)
os.remove(infopath+"/"+fname+".trashinfo")
```

—Rajiv Bhandari, raju.r.bhandari@gmail.com

Simulating the *wc -l* command using *sed*

Myriad tasks can be done with the *sed* command. The one liner below will simulate the *wc -l* command which counts the number of lines present in the file.

```
[mickey]$ sed -n '$ =' file.txt

16
```

Let us demystify the above command. But before that, let's understand the working of the *sed* command. *sed* reads a line from the input file into a pattern buffer (the internal buffer used by *sed*), applies commands (if any) on the pattern buffer, and finally prints the modified line on the standard output stream.

Here, the '=' command prints the line number followed by its contents, '$ =' prints the last line number and its contents, and the '-n' option suppresses the default printing of the pattern buffer. Hence it displays only the last line number. Isn't it interesting?

—Narendra Kangralkar, narendrakangralkar@gmail.com

Tabbed multi-window management and screen sharing using Byobu

Byobu is a GPLv3 open source text based window manager and terminal multiplexer. It's designed to provide elegant enhancements to the otherwise functional, plain and practical GNU Screen, for the Ubuntu server distribution. Byobu now includes an enhanced profile, convenient key bindings, configuration utilities, and toggle-able system status notifications for both the GNU Screen window manager and the more modern Tmux terminal multiplexer. It works on most Linux, BSD and Mac distributions.

To try this, you have to first install *Byobu:*

```
sudo apt-get install byobu
```

—Akhilesh Joshi, joshiakhil.109@gmail.com

Password protection for PDF files

Password protection is a simpler option to secure and transfer files, than encryption with public or private keys. With 'pdftk' you can add, delete and update passwords for a given PDF file. To add a password type, use the following command:

```
pdftk <input_pdf>  output output.pdf user_pw  <password>
```

To delete a password type, use the command given below:

```
pdftk <input_pdf> input_pw <password>  output output.pdf
```

To update a password type, use the following command:

```
pdftk <input_pdf> input_pw <old_password>  output output.
pdf user_pw  <new_password>
```

In all cases, the output file is saved in *output.pdf.*

—Manu Konchady, mkonchady@gmail.com

How to find the Top 10 users of the most space

To display the Top 10 users of the most space, execute the following command at the command prompt/terminal:

```
#du -s /home/* | sort -nr |head -10
```

—Suresh Jagtap, smjagtap@gmail.com

Searching and installing a long list of packages

We often come across the issue of a long list of packages that we need to search:

```
apt-cache search opencv
```

When we complete the search, we get a long list of packages that needs to be copied or installed.

To overcome this issue, we can use the following command:

```
apt-get install $(apt-cache search opencv | awk '{print $1}' |
tr '[:space:]' ' ')
```

This will search and install all packages automatically.

—Rahul Jain, talentediq@gmail.com

Installing GRUB from a rescue CD/DVD

Sometimes it is necessary to reinstall/install GRUB in dual boot systems or single boot systems. To install GRUB:

1. Boot to the live CD/DVD/USB or rescue CD/DVD/USB (nowadays, most Linux distros come with GRUB).
2. Mount the partition that contains your Linux OS, as follows:

```
sudo mount /dev/sda2 /mnt
```

Note: *sda2* is my Linux partition.

3. Install GRUB using the following command:

```
sudo grub-install --root-directory=/mnt /dev/sda
```

Note: The above command is in a single line and *sda* is my HDD.

4. If any errors occur, like you are unable to copy the byte and need to copy the blocks, then just type the following command with *force:*

```
sudo grub-install --root-directory=/mnt /dev/sda --force
```

5. Reboot the system. Now you can see GRUB with OS menus.
6. If yours is a dual-boot Windows/Linux system, then just log into the Linux partition, open the terminal and type the following command:

```
sudo update-grub
```

This will add the Windows OS menu to GRUB. From the next time onwards, you can see the Windows menu on booting up.

—Mahesh Gaali, mahe1729@gmail.com

Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in *OSFY* so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at *www.opensourceforu.com*. The sender of each published tip will get a T-shirt.

DVD OF THE MONTH

Linux for everyone

Ubuntu 15.10 Live Desktop

Ubuntu comes with everything you need to run your organisation, school, home or enterprise. The new Ubuntu release features version 4.2 of the Linux kernel, updated packages of Firefox, LibreOffice and the GNU Compiler Collection along with several bug fixes. The distribution is codenamed `Wily Werewolf' and is one of the most popular Linux distros being used on the desktop. The bundled DVD is a live edition of Ubuntu Desktop for the 64-bit architecture.

The other_isos folder on the root of the DVD also contains an ISO image of Ubuntu 15.10 Server.

Fedora 23 Workstation (Live)

Fedora Workstation is a polished, easy-to-use operating system for laptop and desktop computers, with a complete set of tools for developers and makers of all kinds. You can find the ISO image in the folder *other_isos* on the root of the DVD. The bundled ISO image is for 64-bit live workstations.

Xubuntu 15.10

Xubuntu is an elegant and easy to use operating system. Xubuntu comes with Xfce, which is a stable, light and configurable desktop environment. Try out the DVD to learn more about its features and contents. You can find the ISO image in the folder *other_isos* on the root of the DVD. The bundled ISO image is for 32-bit live systems.

What is a live DVD?

A live CD/DVD or live disk contains a bootable operating system, the core program of any computer, which is designed to run all your other programs and manage your hardware and software.

Live CDs/DVDs have the ability to run a complete, modern OS on a computer even without secondary storage, such as a hard disk drive. The CD/DVD directly runs the OS and other applications from the DVD drive itself. Thus, a live disk allows you to try the OS before you install it without erasing or installing anything on your current system. Such disks are used to demonstrate features or try out a release. They are also used for testing hardware functionality, before actual installation. To run a live DVD, you need to boot your computer using the disk in the ROM drive. To learn how to set a boot device in BIOS, please refer to the hardware documentation for your computer/laptop.

BIG DATA, ANALYTICS & INSIGHTS SUMMIT

4th February 2016

Taj Lands End, Mumbai

Putting Data at the Heart of the Business

At Big Data, Analytics & Insights Summit, we showcase Big Data use cases and techniques that drive the greatest business value. With an emphasis on real-life implementation of Big Data technologies, this practical business forum will provide bold vision from leading innovators across the data-driven spectrum. We explore the latest tools and trends to help delegates gain fresh insights and strategic momentum to grow and respond to their business requirements. Join a select group of professionals, business strategists, BI professionals and data scientists looking to leverage the data opportunity for competitive advantage.

DISCUSSION THEMES

- Predictive Analytics: Staying One Step Ahead of the Connected Customer
- Big Data in the Age of the Customer
- Exploring the Disruptive Potential of Big Data as a Game-Changing Technology
- The Creative Power of Big Data Analytics to Drive New Products, Services and Business Models

3rd February 2016

BIG Data AwardS

Taj Lands End Mumbai

AWARD CATEGORIES

RECOGNISING EXCELLENCE :

- Big Data Start-up
- Big Data Analytics
- Big Data Tech Provider
- Data Visualization
- Business Intelligence
- Data Governance
- Data Security
- Open Source Tool
- Industry Application Award

MOST ADMIRED :

- Big Data Start Up
- Big Data Project
- Big Data Tech Provider
- Data Warehouse of the Year
- Architectures in Operational Analytics
- Architectures in Security & Compliance

BEST DATA PRACTICES :

- Data Warehousing
- Data Mining Strategy
- Data Collection Strategy
- Data Analysis Strategy
- Emerging Technologies

BEST ANALYTICS SERVICE IN :

- Business Intelligence
- Marketing Analytics
- Predictive Modelling
- Sales Analytics
- Workforce Analytics

INDIVIDUAL AWARDS :

- Business Analyst
- Marketing Technologist
- Analytics and Insight Leader
- Data Security Programme Leader
- Data Entrepreneur / Founder

CRÈME-DE-LE-CRÈME :

- Innovative Technology of the Year

www.thecustomerfest.com

For Further Details Contact

Aboli : +91 9833226990 / +91 9004958990

harish@kamikaze.co.in / aboli@kamikaze.co.in

Media Partner

OpenSource For You

EFYGROUP

Conceived & Managed by

KamiKaze

Kshitij, 103 – 1st Floor, Veera Desai Road, Opp Andheri Sports Complex, Andheri (W), Mumbai 400 058
+91 22 613 81800 / 26780476

R N I No. DELENG/2012/49440, Mailed on 27/28th of Advance month Delhi Postal Regd. No. DL(S)-01/3443/2013-15
Published on 27th of Advance month Licenced to Post without Pre-Payment Licence No. U(S)-56/2015